AUSTRALIAN
communications
JULY 1993 The Networking and Telecommunications Management Magazine $5.50

THE SNA MIGRATION DILEMMA
How to position the network for LAN interconnect?

RATING ROUTERS
Can routers be trusted with critical data?

LAN MANAGEMENT
Will new proactive tools keep LANs trouble-free?

INTEL'S PENTIUM
What does the new processor mean for net managers?

CDMA
Is it the future for mobile communications?

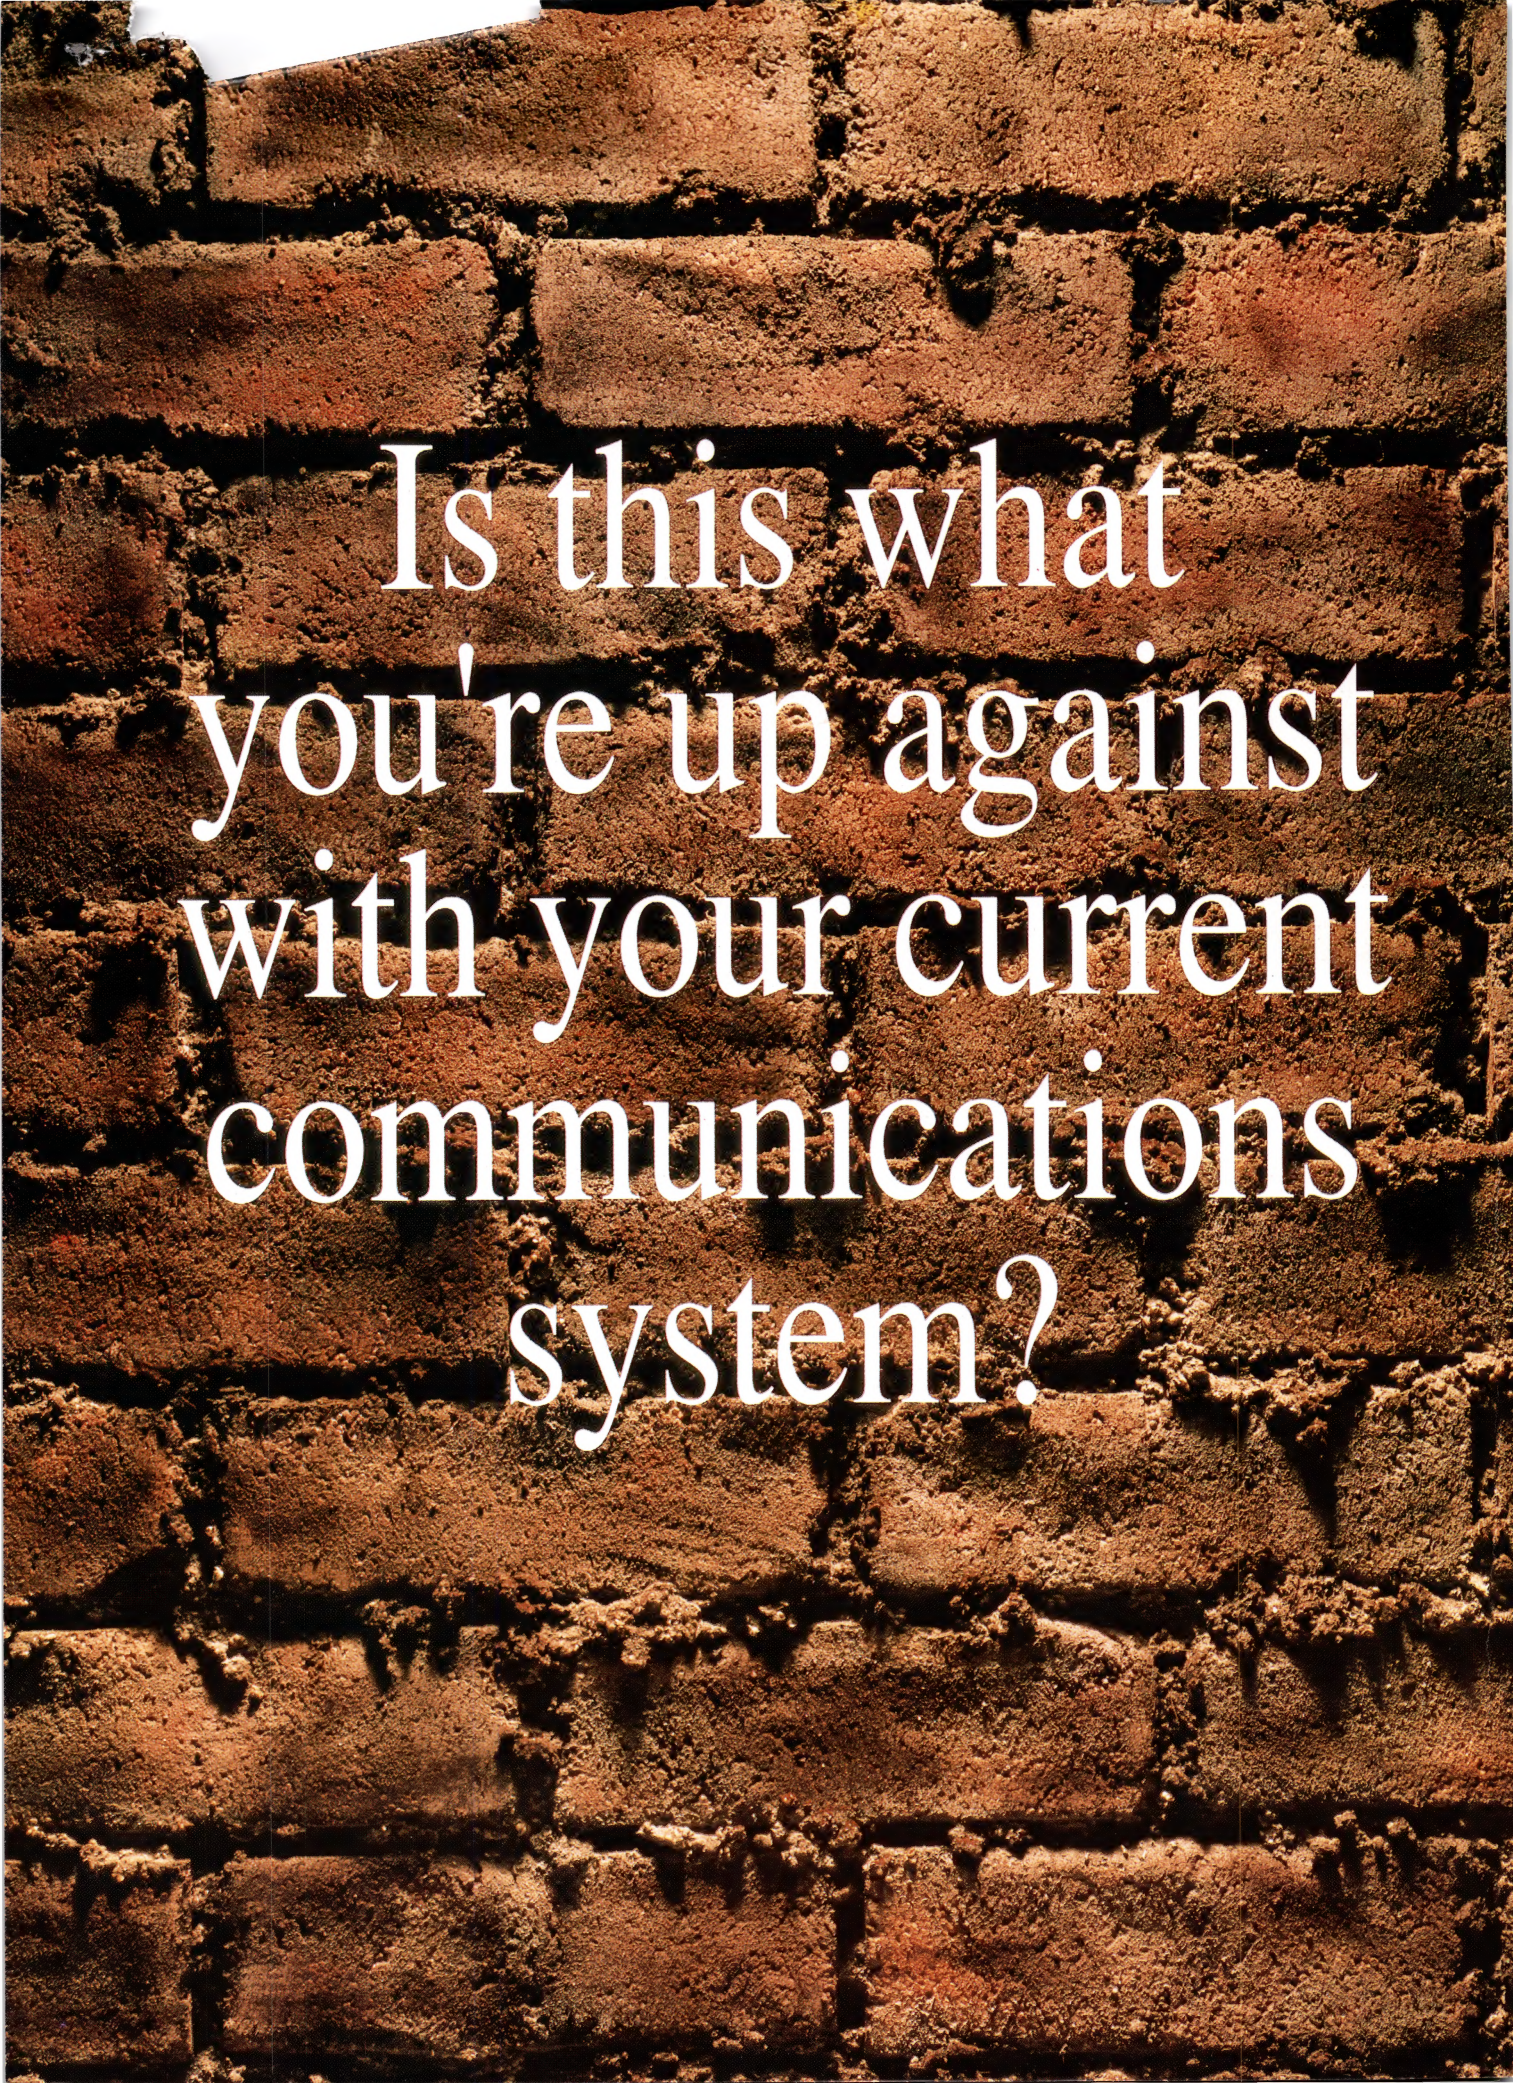
Is this what
you're up against
with your current
communications
system?

AT&T opens up endless possibilities with the DEFINITY Communications System.

Face it. You've spent hundreds of thousands of dollars to bring your company's phone system into the eighties. Only to discover it can't make it into the nineties. Do You: a) Admit you're only human. b) Try to squeeze another six months out of your current system. c) Replace your current system with the latest model that's built for obsolescence. d) Invest in the last communications system you'll ever need.

If you answered 'd,' your choice is AT&T's DEFINITY Communications System. That's because the DEFINITY Communications System is the only system with a modular configuration that takes you from 40 lines to a seemingly endless 30,000.

In most cases, you can add those extra lines, simply by adding on to the modular base. And AT&T the world's largest Telecommunications Company now in Australia, can help you plan those additions, so you don't have to shut out your customers' calls.

So if your old PABX has brought you to a dead end, the DEFINITY Communications System can turn your fortunes around.

To find out more about the DEFINITY Communications System, contact your AT&T Account Executive, or call (02) 352 9000 Fax (02) 352 9111.

For PABX technology, it's a breakthrough.

AT&T
The right choice.

AT&T BUSINESS COMMUNICATIONS SYSTEMS AUSTRALIA, 4A LORD STREET, BOTANY 2019.

AUSTRALIAN
communications

Cover: John Dymond Photography

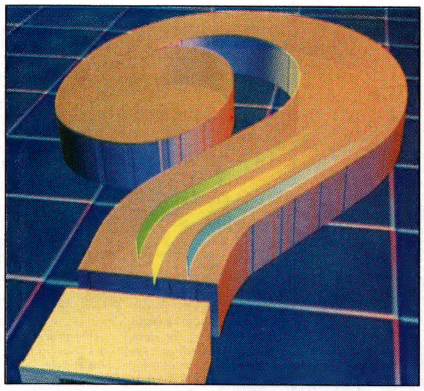

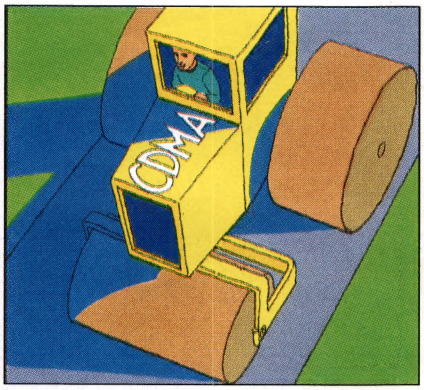

JULY 1993

LAN MANAGEMENT 89

Behind-the-scenes services that keep applications running smoothly are taken for granted in the mainframe world, but they are spotty at best on LANs. LAN management packages must go beyond system troubleshooting to spot problems before they occur — in other words, they must become proactive. Makers of LAN monitors, protocol analysers, management systems, smart hubs and simulation tools are now adding to their products the intelligence needed for proactive management. For proactive management to become more than a buzzword, however, all these different products and approaches must be able to work together.

Architectural Issues

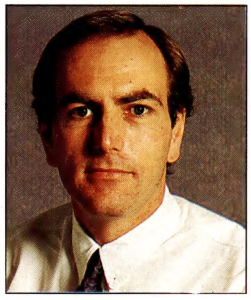

Communications architectures of vastly different types feature strongly in this month's edition. In our cover story, Peter Johnson, a leading Australian IBM Systems Network Architecture expert takes a look at the available and emerging options for SNA network managers. The challenge they now face is to evolve their legacy networks to successfully handle emerging LAN interconnect requirements.

According to Peter, the fundamental choice is between evolving the hierarchical SNA network to IBM's Advanced Peer-to-Peer Networking (APPN) or building an enterprise multiprotocol router network. In our June edition (see 'APPN Rises to the Enterprise Challenge') we examined APPN and perhaps its chief architectural antagonist, TCP/IP, at length. Peter builds on this platform by taking a look at some APPN migration strategies and some products available now which transparently connect non-SNA LAN traffic across SNA WANs. He also provides an update on Advanced Peer-to-Peer Internetworking (APPI), a Cisco Systems-inspired initiative designed to provide an open, TCP/IP-oriented solution for SNA peer-to-peer networking.

Of course the feasibility of moving to a multiprotocol router network is dependent to a great extent on the ability of routers to cope with mission critical network traffic, so our review of router prioritisation schemes and performance ('Can Routers Be Trusted With Critical Data?' starting on page 73) will make interesting reading for network managers now facing the SNA dilemma.

Mobile communications architectures are also examined closely this month as Stewart Fist takes a look at GSM and CDMA mobile technology ('Grand Scale Mistake?' on page 13 and 'Will GSM and D-AMPS Give Way to the CDMA Push?' on page 83). Selected by Austel as the most suitable first digital mobile communications system for Australia, GSM's apparent flaws are now being widely reported on around the world. Like D-AMPS, its US digital mobile cousin, GSM employs TDMA access techniques which some blame for problems with hearing-aids and other electrical devices. If the problems are real and — even worse — virtually impossible to fix, as some would contend, then GSM would appear to be in serious trouble — as would a few of its European backers. Short of subscribing to politically correct theories about a conspiracy of silence imposed by the powerful multinational interests behind the technology, it is a little difficult to believe that GSM could be fatally flawed. More plausible is the line that it indeed does have problems, but ones that can be overcome — albeit not easily.

But perhaps a bigger challenge than real or over-hyped technical problems for GSM is the CDMA technology now being refined by US firm Qualcomm. Stewart has reported at length in a previous issue (see 'GSM — Why Are We Rushing to Embrace it?' *Australian Communications*, February 1992) on Qualcomm's technology, and this month he provides an update on its status and success in attracting interest from the US RBOCs.

MA Smeaton

AUSTRALIAN communications

EDITOR
Mark Smeaton

CONTRIBUTORS
Tom Amos, Stewart Fist, Liz Fell, Peter Leonard, Peter Waters, Stuart Corner, Graeme Le Roux, Bill Bennett, Bernard Levy, Shelley Spriggs, Robert Clark

ASSISTANT EDITOR
Sarah Parkes

PRODUCTION MANAGER
Samantha Keats

ILLUSTRATIONS
Samantha Keats, Louis Silvestro

SALES AND MARKETING DIRECTOR
Craig Burkill

CHIEF FINANCIAL OFFICER
Jennifer Stibbard

Subscriptions
Australian Communications is available by paid subscription for $54 per annum. Subscribers receive eleven issues per annum.

Overseas Rates by Air	1 Year	2 Years
New Zealand, Papua New Guinea	$66	$110
Singapore, Indonesia, Malaysia, Brunei, Pacific Islands	$72	$122
Asia — including Hong Kong, India, Korea, Japan, Taiwan	$78	$133
Europe, North America, Middle East, South America	$92	$168

Publishers
Published and distributed monthly by Ostasun Pty. Ltd. (ACN 003 606 102) under licence from ACP Computer Publications, a division of ACP Publishing Pty. Ltd. (ACN 053 273 546) of 54 Park Street, Sydney NSW 2000. ISSN 0818-9021. **Address:** Level 4, 541 Kent Street, Sydney NSW 2000. **Tel:** (02) 264 2200 **Fax:** (02) 264 2244.

Printed at **Offset Alpine Printing**, Derby & Wetherill Streets, Silverwater NSW 2141

Member of the Audit Bureau of Circulations

abc

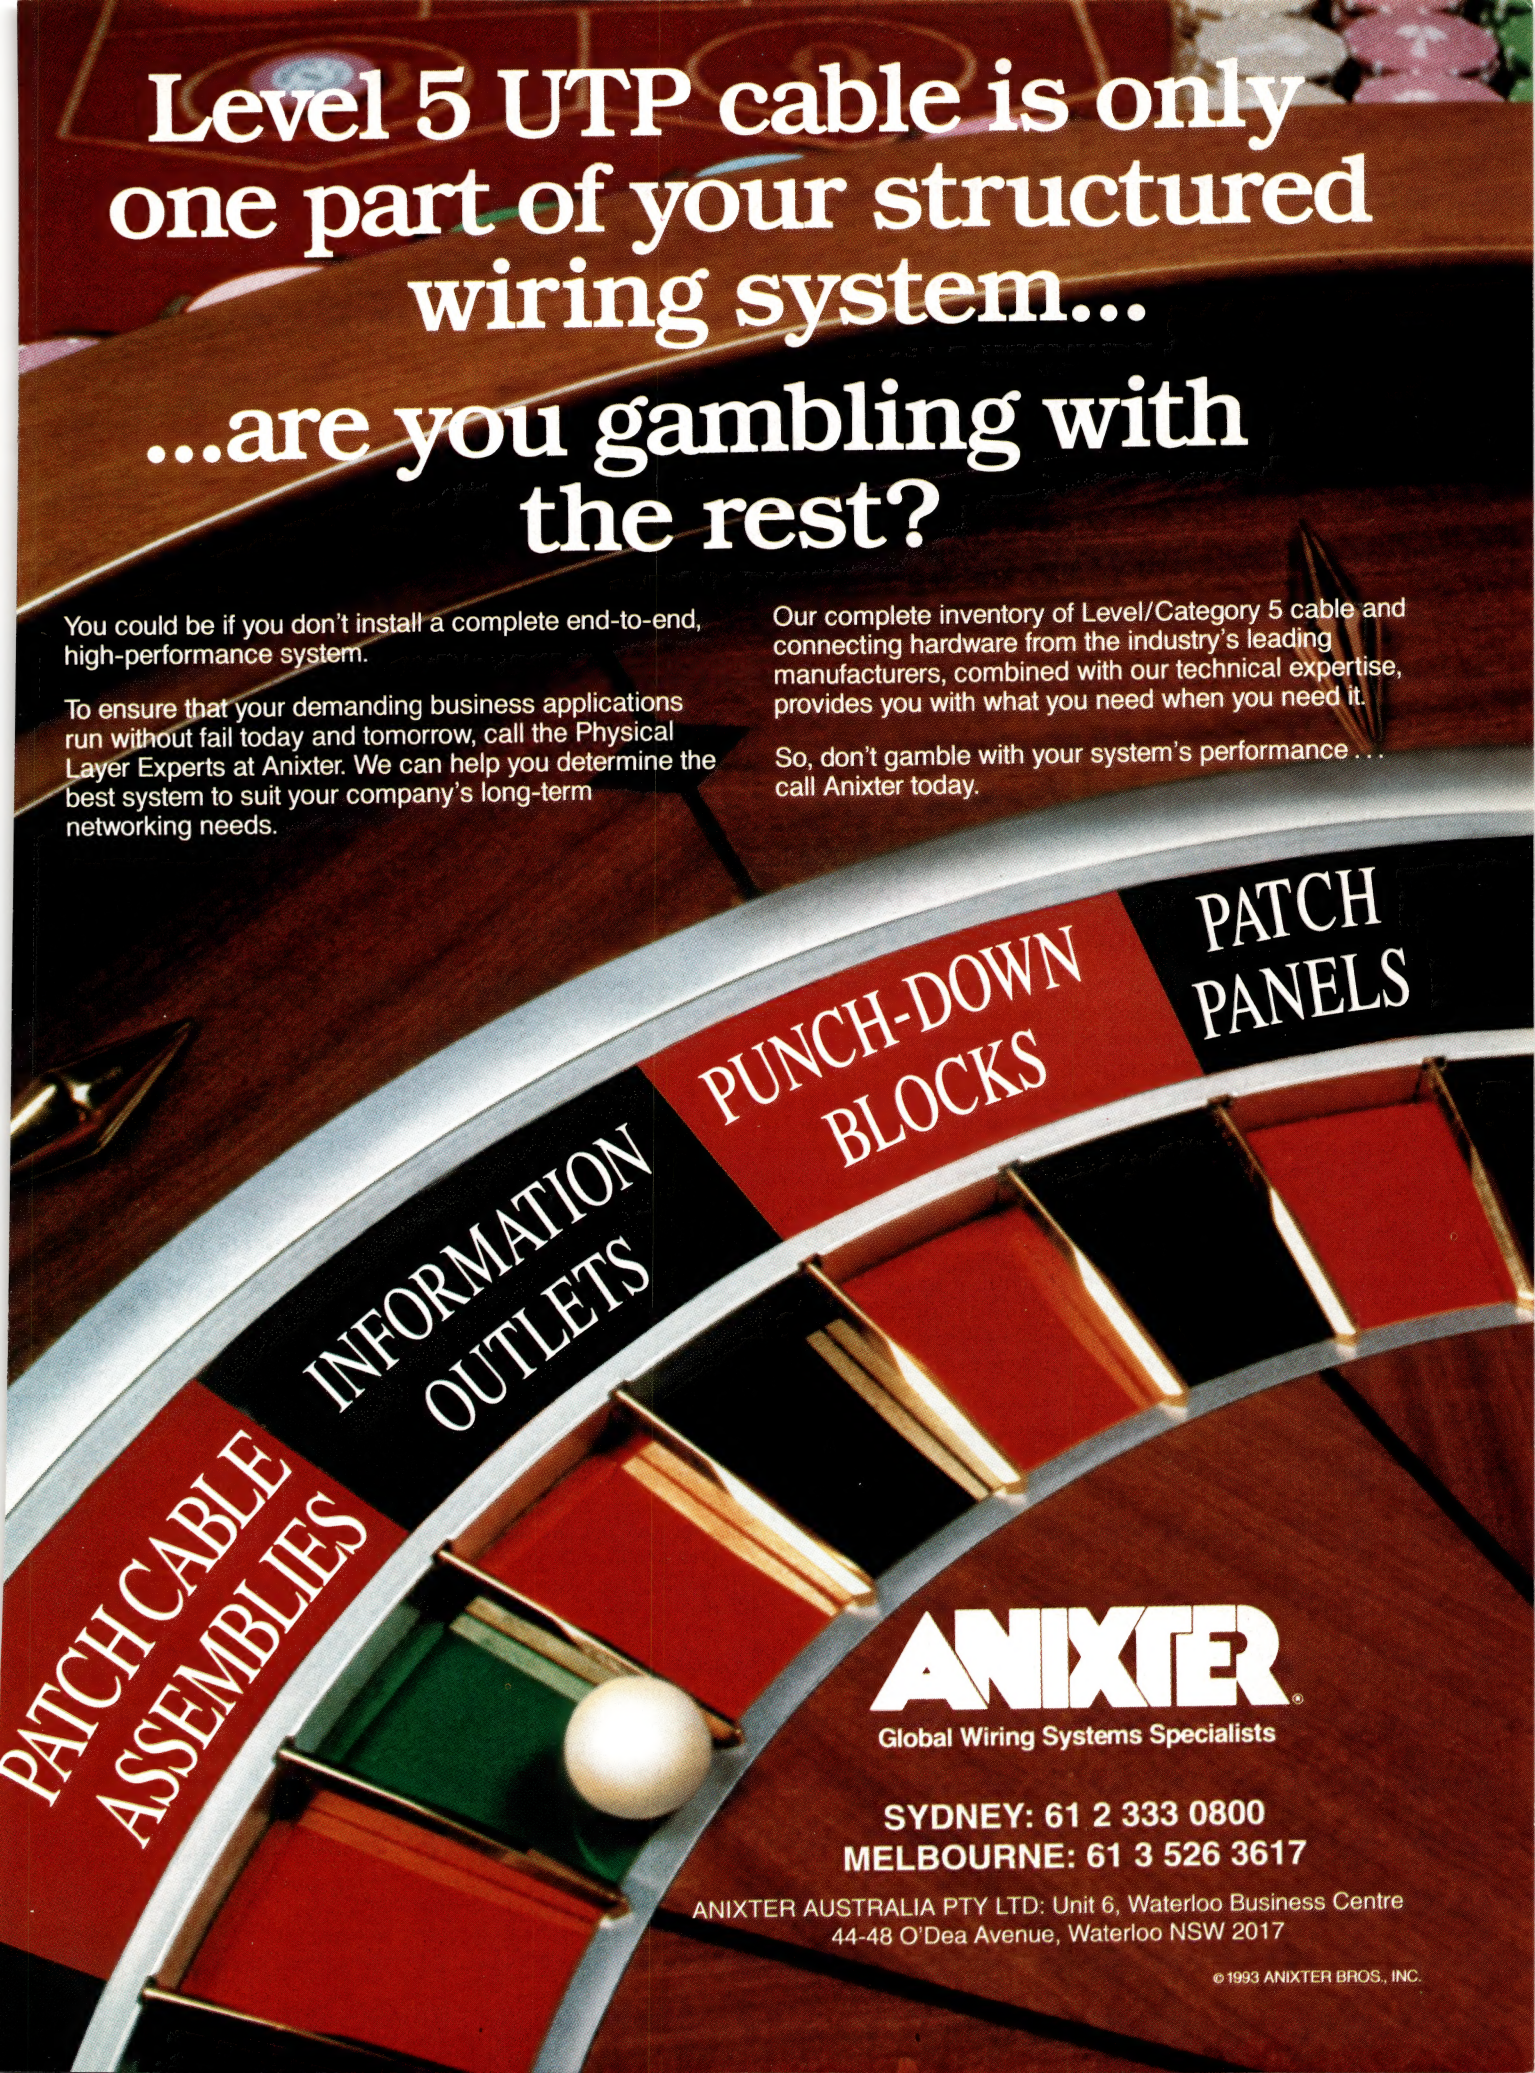
Level 5 UTP cable is only one part of your structured wiring system...
...are you gambling with the rest?
You could be if you don't install a complete end-to-end, high-performance system.
To ensure that your demanding business applications run without fail today and tomorrow, call the Physical Layer Experts at Anixter. We can help you determine the best system to suit your company's long-term networking needs.
Our complete inventory of Level/Category 5 cable and connecting hardware from the industry's leading manufacturers, combined with our technical expertise, provides you with what you need when you need it.
So, don't gamble with your system's performance . . . call Anixter today.
PATCH PANELS
PUNCH-DOWN BLOCKS
INFORMATION OUTLETS
PATCH CABLE ASSEMBLIES
ANIXTER
Global Wiring Systems Specialists
SYDNEY: 61 2 333 0800
MELBOURNE: 61 3 526 3617
ANIXTER AUSTRALIA PTY LTD: Unit 6, Waterloo Business Centre
44-48 O'Dea Avenue, Waterloo NSW 2017
© 1993 ANIXTER BROS., INC.

Let the Preselection Battle Begin

Telecom and Optus last month began preparations for the first skirmish of their preselection ballot war. Commencing July 15, telephone customers in Canberra will receive ballot forms through the mail requesting they nominate their carrier of choice. Those not responding within 30 days will remain with Telecom.

Prior to Optus' selection of Canberra, Austel was forced to arbitrate on the 'scope, thrust and cost' of its public ballot education campaign after disagreement between the carriers.

The industry regulator determined early last month that the cost of the campaign will be approximately $5 million (jointly funded by the carriers), and will involve TV, radio, print media, direct mail and 008 information lines. Most importantly, Austel ruled that the campaign will not focus on ballot response rates and that encouraging participation should be left to the carriers.

The dispute over the public education campaign coincided with a reprimand from the Trade Practices Commission for both carriers for misleading advertising, with Optus agreeing to withdraw some offending ads, and Telecom being forced to run full-page corrective advertisements in leading newspapers.

Telecom and Optus did manage to agree on a ballot administrator and have chosen a consortium led by Price Waterhouse Urwick. The administrator will distribute and collect ballot papers and ensure each respondents' choice is implemented.

IDA Report Shows Exports Rising

Continuing growth in the cellular mobiles market has helped boost the combined turnover of companies participating in the Industry Development Arrangements (IDAs), Austel's most recent report shows.

The report, which covers the first nine months of the fourth year of the IDAs, shows that turnover reached $299.36 million for the nine month period ended March 1993, compared to $262.70 million for the corresponding period last year.

Local content fell by 8.3% against turnover, and also down 12.5% against Australian production when compared to the first nine months of the first year of the IDAs, and down 2.6% against turnover and around 3.9% against Australian production when compared with the same period for the third year.

However, exports continued to rise, increasing by $26.69 million to reach $60.72 million, or 20% of turnover. Imports also rose to reach 52% of turnover.

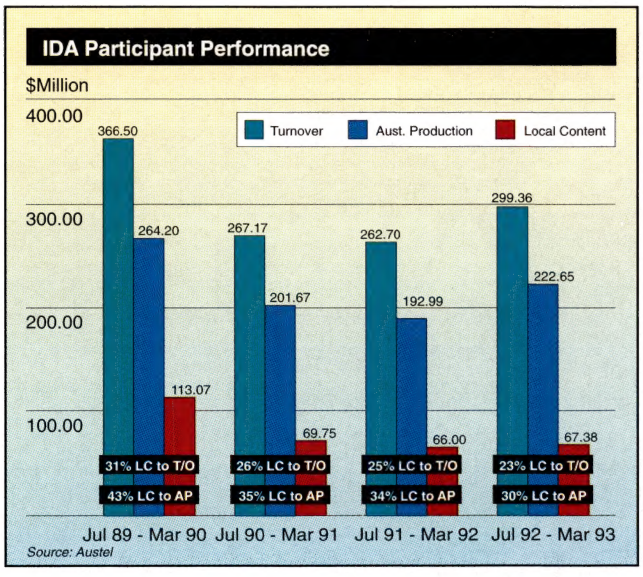

Optus Extends Satellite Services

Optus has announced a large scale re-organisation of its satellite services and pricing. The changes took effect on July 1, and include new transmission options and lower prices for existing services. A spokesperson said the company will offer a new range of service options closely aligned with the needs of the broadcasting industry.

Optus' Australian domestic satellite network is comprised of three satellites and a network of earth stations across the country. It is used for TV, telephone calls, radio, civil aviation signals and a wide range of other services.

The spokesperson said the new services range from new transponder capacity and extensions to the basic service, to a proposed virtual private network, which the company plans to have available to customers within two years.

OPTUS SATELLITE TRANSPONDER CHARGES

		92/93 Price (1 trans)	93/94 Price (2nd trans)	93/94 Price (1st trans)
Annual Lease				
National Beam	Level 1	$7.800M	$7.185M	$7.332M
	Level 2	$6.800M	$6.265M	$6.392M
Spot Beam	Level 1	$6.256M	$5.763M	$5.881M
	Level 2	$5.441M	$5.011M	$5.115M
Fixed Term Lease (5 years)				
National Beam	Level 1		$6.082M	$6.941M
	Level 2		$5.931M	$6.052M
Spot Beam	Level 1		$5.456M	$5.567M
	Level 2		$4.744M	$4.841M
Fixed Term Lease (10 years)				
National Beam	Level 1		$6.424M	$6.555M
	Level 2		$5.601M	$5.715M
Spot Beam	Level 1		$5.152M	$5.257M
	Level 2		$4.480M	$4.572M

Austel Billing Report Released

Austel has released its report on Call Charging and Billing. The main outcome of the report is that carriers will be required to provide reports on the accuracy of their billing systems. Austel also plans to set targets for charging accuracy, and says carriers failing to meet these targets will be deemed to have breached their licence conditions.

The report makes 19 recommendations but fails to fully endorse many of Austel's original proposals, including the installation of 'telemeters' in customer premises; the provision of full details of all calls, including the routing of long-distance calls, on customers' accounts; and the introduction of comparative reporting methods, similar to those used for some electricity, gas and water bills.

Telstra to Join AT&T Alliance

AT&T has announced a new alliance designed to offer global services to multinational customers. Called WorldPartners, the alliance initially comprised AT&T, Japan's Kokusai Denshin Denwa and Singapore Telecom, with Telstra, Korea Telecom and Canada's Unitel to join soon, and other PTTs expected to follow.

The WorldPartners alliance will provide a new family of business communications services known as WorldSource. The WorldSource service will be launched in the US and the Asia-Pacific region by the end of the year. It will provide customers needing global communications links to overseas offices with virtual private networking facilities, private voice/data circuits and frame relay.

Optus COO, Ian Boatman, (third from left) with Krone staff members

Krone Wins Optus Contract

Krone Australia has won what it terms a 'significant' contract to supply Optus with Australian manufactured network termination equipment.

Under the terms of the agreement, Krone will supply central office and premises wiring systems to Optus' strategic suppliers, Fujitsu and Nortel. Krone LSA Plus copper termination equipment will also be installed in telephone exchanges in each State. The equipment will provide the point of termination between the Optus network and the customer's in-office wiring system.

An Optus spokesperson said Krone had been chosen because the company was an Australian manufacturer able to help Optus increase the Australian content of its service, adding that Krone has Australian standards certification, is a tried and tested supplier which has had a long relationship with Aussat, and has a reputation for high quality products and on-time delivery.

ITU Predicts Asian Telecom Boom

A study by the International Telecommunications Union predicts high growth rates for telecommunication development in the Asia-Pacific region. The report, entitled *Asia-Pacific Telecommunication Indicators*, reviews the telecommunications networks of 39 countries in the region, based on data up to 1991.

It shows that the Asia-Pacific region is the world's fastest growing, with growth rates considerably higher than those of Europe or the Americas. Over the last decade the region has experienced steady growth of 7% per year, leading to an increase of 70% in the number of main lines. Main line growth in the region's faster growing economies of Hong Kong, Singapore, and the Republic of Korea has been around 14%.

The ITU report also reveals that exchange line digitisation in Hong Kong and Korea is around 90% — much higher than that of Australia, which was only 26% in 1991, despite a higher teledensity (measured as the number of main lines per 100 inhabitants).

The report also confirms the importance of growing trading links within the region. Within-region calls accounted for about 55% of all traffic originating in the Asia-Pacific region — higher than the North American Free Trade Area and around the same as the European Community, despite the much larger distances involved. China, in particular, has experienced very strong growth in international traffic, in the order of 50% per year.

In contrast, Japanese carriers NTT and KDD have experienced depressed revenue growth.

The market for mobile telephones is predicted to boom, particularly in countries with poor fixed-link infrastructures. Growth rates of 90% p.a. compare with annual growth of fixed systems of just 7%. The report says mobile phones are likely to outnumber fixed link phones by 1995 in the Philippines, and 1997 in Thailand.

In Brief

Alcatel Alsthom SA has announced a $500 million Commercial Paper Issue for Australia, which a company spokesperson said was recognition of the significance and growth potential of the Australian market. The issue will provide working capital finance for Alcatel Australia and its subsidiary GEC Alsthom.

Microsoft has announced it will improve its support services with the introduction of a new support strategy. The new service will comprise three support levels — Free Support, Express Support, which involves a 'pay-for' service, and Premier Support, which is designed for large customers and channel partners.

JNA has obtained exclusive marketing rights to all Design 2000's ISDN products and developments in a new five-year agreement. JNA said the agreement would mean its customers will be able to reduce communications costs for retail operations such as EFTPOS, and those involving links with small remote facilities.

PacRim Financial Network has launched its FlexiFax local and international broadcast facsimile service. The new service allows users to transmit fax documents from standalone PC or Macintosh computers or from a Novell LAN, to a client's fax machine using PacRim's existing value-added network.

SynOptics has announced the launch of its SynOptics Affiliate Reseller (SAR) program, which the company says is designed to build relationships and expertise in SynOptics' networking solutions among specially selected dealers and resellers. The program currently involves five resellers.

The NSW Department of Health has awarded a contract to develop a large strategic telecommunications plan to a consortium involving Gibson Quai & Associates and Ernst & Young. The plan covers the whole of the NSW Public Health System for the next five years, and addresses voice, data and imaging requirements.

Scitec is to join Telecom's Enhancement Distribution Channel. The company will offer Telecom services such as ISDN lines, Megalink and DDS, bundled with its Maxima series of networking products, and Telecom will be able to incorporate Scitec products into telecommunications packages for its customers.

Pacific Star has entered into a partnership agreement with the Queensland Government, whereby 15 technical staff from the Centre for Information Technology and Communications will work from Pacific Star's offices to assist in managing the Government's $100 million-a-year telecommunications network.

DEC has just completed one of its Australian biggest systems integration projects. The $16 million system for BP Australia is known as Petrolink, and connects all of BP's distributors across the country via a single network.

Com Tech Communications has been appointed a distributor for Eicon Technology, a world leader in the connection of PCs and LANs to X.25 packet-switched networks.

Olex Cables has just opened a third facility in China, bringing its total investment in that country to $50 million. A new Beijing office will manage current and future projects in China, and supplements two factories in Shenzen and Tianjin.

Wang and Telecom Australia have signed an agreement to promote a closer selling relationship between their sales forces to increase joint business opportunities. The arrangement is part of Telecom's Enhanced Distribution Channel Program.

NetArch has installed what it claims is Australia's largest secure fibre optic network for the Department of Defence in Canberra. Company officials claim the network is the first totally secure system in the world combining Ethernet, FDDI, 3270 and synchronous and asynchronous protocols over optical fibre.

Sakhalin Island has officially opened its first satellite links with the rest of the world via a new satellite earth station in its capital, Yuhzo-Sakhalinsk. Installed by Telstra, the standard 7.5 metre F2 earth station will use the Intelsat system to link the Far Eastern Russian province to Australia and the rest of the world.

Gerry Moriarty, formerly Managing Director of BCL, the primary network supplier to NZ's Clear Communications, has joined Telecom Australia as Managing Director of Network Products.

In Brief

Cisco Systems has reported net income of $US46.3 million and net sales of $US172.3 million for the third quarter ended April 25, 1993. The figures represent increases of 99% and 89% respectively over the same period last year.

GPT and Ericsson have won contracts worth £12 million and £17 million to supply Synchronous Digital Hierarchy (SDH) transmission to BT. GPT equipment will be used to link academic MANs as part of the SuperJanet academic network, while Ericsson equipment will be used to send traffic from cable landing points and satellite earth stations to major cities.

Pacific Transit Cable group formed by the Pacific Telecommunications Council is to look at the feasibility of building an undersea fibre cable linking South America with New Zealand. The cable could start in Chile and interconnect with Australia.

BT and Mercury will soon have some competition from satellite operators. The UK Department of Trade and Industry has licenced Satellite Information Services, E-Sat and Panamsat to provide a wide range of international services by satellite.

Alcatel, Nokia and Ericsson have all announced contracts to supply Thai fixed or mobile operators. Alcatel has won a $US60-70 million contract, Nokia has signed an agreement worth FIM 150 million, and Ericsson has won a $US130 million contract.

Telstra, C&W and Benpres have signed a Memorandum of Understanding to carry out a feasibility study on the installation of 1 million new phone lines in the Philippines.

Poland's cellular system now covers 20% of the country's population. The network, installed by Centretel, a Polish PTT, Ameritech and France Telecom, will cover all the main trunk roads and corridors between the nine Polish cities by the end of 1993.

Alcatel is to supply a nationwide Metropolitan Area Network to the Austrian PTT. Around 150 34Mbps and 500 2Mbps subscriber access ports are due to be installed by the company over the next two years to build the network.

Ericsson has announced it is extending its five year old cooperation with the Beijing Wire and Communications Plant (BWCP) China in the production of PABX systems.

Northern Telecom Europe has won a second contract to supply transmission systems to Egyptian operator, Arento. Northern will supply a 140Mbps optical line transmission system to form the backbone of a network serving the Fayoum area of Upper Egypt.

BT and EDS have backed away from a proposed alliance. The alliance would have involved a quarter of GM class E shares to BT, which would have represented a claim on 25% of EDS profits but not an ownership stake in the company.

Ericsson's consolidated net sales in public telecommunications, radio communications and business networks rose by 36% to SEK12,867 million in the first quarter of 1993, with radiocommunications showing a particularly strong increase of 54%.

Northern Telecom has announced that it has won contracts to deliver Synchronous Digital Hierarchy (SDH) equipment to two Russian operators, Rascom and Metrocom.

International Digital Communications (IDC) of Japan has joined Cable and Wireless, US Sprint and five other carriers to provide a global virtual private network (GVPN) service in Japan.

GPT Video Systems has announced a contract to supply China's first videoconferencing network. The network will be operated by Hunan Post and Telecommunications.

Novell has reported revenues of $US281 million for its second quarter ended May 1 1993, up 25% from revenues for the same period last year. Net income increased 31% to $US80 million.

Inmarsat and DB Salyut have signed a contract for the first launch of a Western-built satellite on a Russian Federation launcher. The $US36 million contract is for the launch of an Inmarsat-3 satellite on a Proton rocket from the Baikonur Cosmodrome, scheduled for 1995. Inmarsat has also elected a new Chairman. He is Jean-Paul Brillaud, who is currently head of the Satellite Communications Policy Division at France Telecom.

MCI, BT Form Strategic Alliance

MCI Communications Corporation has announced the formation of a strategic alliance with BT in a move designed to enable the two companies to lock horns with the giant AT&T in a wrestle for a bigger share of the global telecommunications market.

The agreement involves the setting up of a new joint venture company, and calls for BT to acquire 20% of MCI for approximately $US4.3 billion in cash, and for MCI to acquire BT North America. Three BT representatives will take places on MCI's Board of Directors, while MCI Chairman and CEO, Bert Roberts, will join the BT Board.

The announcement has been interpreted by industry observers as a major assault on AT&T. The rivalry between BT and AT&T has grown increasingly acrimonious over the past few months, as both carriers compete aggressively for a share of the global market. MCI, which now ranks as the second biggest long-distance communications carrier in the US, has also been forming alliances with a number of foreign carriers as a means of funding head-on competition with AT&T in all markets.

The new partnership involves the establishment of a joint venture company which will provide seamless value added voice and data services to customers throughout the world. BT and MCI are expected to invest over $US1 billion in the company, with BT holding approximately 75% of the equity, and MCI the remaining amount.

MCI said the two companies will operate the venture as equal partners, and will both market new jointly branded products and services created by the new firm. MCI will be responsible for North, South and Central America and the Caribbean, and BT will handle the rest of the world. Traffic from these new services will be transported via a new global intelligent network which will be established by the joint venture.

BT Profits Drop By £1 Billion

BT has suffered a decline in profits of over £1 billion to £1,972 million in the year to March 31. Turnover was down 0.7% to £13,242 million, despite growth of 2.8% in turnover in the fourth quarter 1992.

Profits were affected in part by efforts to make up a £750 million deficit as of January 1st 1993 in the company's merged pension scheme. A BT official said executive directors' salaries will be frozen until a scheduled review in July 1994.

Other figures show costs fell 1.3% during the year before allowing for staff redundancy payments. Almost 40,000 employees left the company in 1992-93, bringing the total number employed at March 31 to 170,000. A further 15,000 will leave the company over the next two financial years. BT estimates the cost of voluntary redundancies in the current financial year will be around £500 million.

NTT to Build Fibre Network

Japan's NTT Corporation plans to invest ¥45 trillion ($A620 billion) in new plant and equipment. The money will be used to develop new services and build a fibre optic network throughout Japan by the year 2015.

NTT President, Mr Masashi Kojima, said that continued capital investment at the rate of ¥2 trillion ($27 billion) each year will enable NTT to develop the network, and will cover costs such as laying fibre subscriber lines and installation of digital switches and carrier equipment.

By 2015, NTT aims to provide VI&P (visual, intelligent and personal) services to every household in Japan. The company will build the network first in cities such as Tokyo, Osaka and Nagoya, and then steadily extend it to other areas.

Mr Kojima indicated that NTT will be seeking to revise its telephone call rates to cover the cost of the investment.

HALF-PRICE ETHERNET*

Ungermann-Bass attacks the cost of managed enterprise networks

Gives you the flexibility to either assign each group of eight ports to one of five backplanes, or remain independent

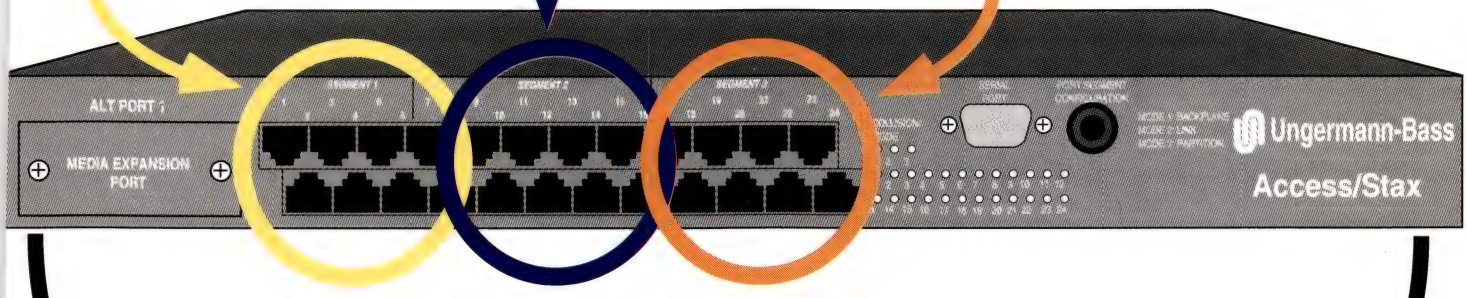

lexible support or non-UTP nedia providing ost-effective ptions for urrent and uture network equirements

Optional Media Connectors

DB-15 connector
AUI (Thicknet)

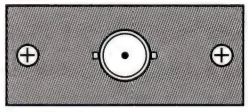

BNC connector
for Thinnet

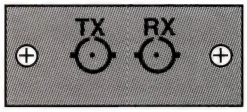

Fibre (FOIRL)
connector

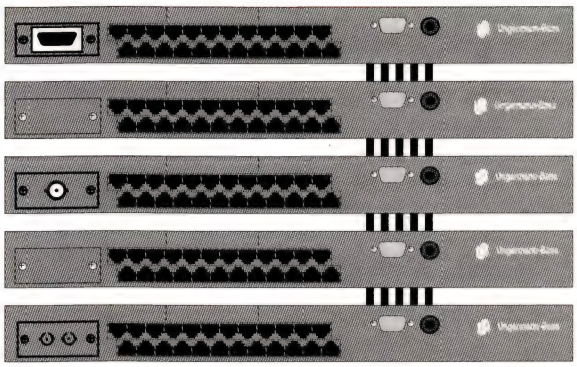

Expandable to five hub stacks, supporting up to 15 Ethernet segments and 120 ports

Flexplane acts as five backplanes plus management bus

SNMP MANAGED ACCESS/STAX FOR AROUND $150 PER PORT **

*Comparing Access/Stax to price of three leading managed Ethernet hubs
**Offer valid until 30th September, 1993.
Based on 120 port bundle ex-tax

BUILT-IN MANAGEMENT

Cost effective real-time management.
SNMP management to 802.3k standard

NETWARE OPTIMISATION

Fully compatible with Novell's NetWare Hub Services (HUBCON) and Hub Services Manager (NMS/HMI)

INTEGRATION

Fully integrates with Ungermann-Bass' enterprise network management architecture

For further information contact Ungermann-Bass on 008 814 499

Some of the world's finest are home grown.

Like Australia's world renowned wines our LAN hardware is second to none!

With a range of world class connectivity products from Australia's only LAN manufacturing company, INC Manufacturing.

All with a 2 year warranty and delivery from stock, INC products are available nationally through selected dealers and integrators.

INC offer a unique range of solutions for

Token Ring

Passive and Active MAU's
Copper Repeaters
Fibre Optic Repeaters

AS400

Passive LWC's
Active Wiring Concentrators
Active Stars
5250 Emulation Boards
5250 Multiplexers

Ethernet

10 Base/T Transceivers
8/16 Mbps Cards
6 and 12 Port Hubs

Call us on 008 807 982.

INC MANUFACTURING

INC Manufacturing Pty Limited A.C.N. 051 977 701
29 Attunga Road, Blaxland NSW Australia
PO Box 182 Glenbrook NSW 2773 Australia
Phone 047 39 6111 Fax 047 39 4247

INC1012

Grand Scale Mistake?

GSM has problems which are now attracting attention around the world. Will action be taken to resolve them?

You'd think that Australia would have learned from past experience that buying high-tech gear straight off the drawing board is fraught with dangers. But, if our choice of GSM is any guide, it seems we haven't. It's easy to criticise after the event, but it is necessary to analyse what went wrong in the way we selected, and then maintained, our stance on GSM, as evidence of problems grew.

First of all, it is elementary physics that, if you turn a transmitter on and off very rapidly, as you must do in TDMA (time division multiple access) systems like GSM, then you are going to generate a low-frequency component (and a couple of harmonics) at the rate of switching. GSM zaps its transmitter 217 times a second, and it generates ELF audio frequency interference at 217Hz, 434Hz and 651Hz. The lowest harmonic has the most power, but the second and third affect hearing-aids the most, since the aids tend to filter out the bass frequencies.

The UK Department of Trade and Industry discovered these emissions at least 18 months ago, and advised ETSI — which, apparently, knew all about it well before. There's a copy of a September 1991 letter circulating from ETSI's Radio Equipment and Systems group which says 'it is a tremendously difficult task' to develop shielding for TDMA equipment. That's putting it mildly — it is actually impossible.

If you want further proof, look at the GSM specifications. The designers provided a set 'ramp-up period' rather than just turning the transmitter on — and the only reason for going to this trouble is to reduce any higher-frequency audio components generated by this on-off action. So there's no doubt that ETSI was well aware of potential problems here years ago. But the attitude seems to have been that, (in the words of one divisional head in DOTAC): "It's a hearing-aid problem, not a cellular phone problem!" And here, they are quite wrong.

I can see a number of side issues arising from this kind of low-frequency emission:

- Intrusions of audio frequencies into other analogue electronic equipment like radio sets, amplifiers, tape recorders, etc.;
- Intrusions of base-band digital pulses into digital equipment; and
- Potential and perceived (even if irrational) health problems.

Shielding hearing-aids with metallic-plastic sprays will not solve these secondary defects in GSM. And they are both real and substantial. One European laboratory claims that GSM produces two-to-three times the emission levels of the accepted European CISPR standards, and Telecom's Research Labs found low-frequency emissions at eight times these levels.

So why do the Europeans bother having emission standards at all?

We must also remember that when a number of GSM handsets are sharing a channel (or even operating in nearby channels) they are pulsing in an eight-bit (once for each slot) pattern. And since the synchronised pattern of slot-use varies, and each transmitter can hop between channels, the cumulative effect must be the random generation of digital baseband pseudo-bytes at a rate of 217 bytes per second — which could also be interpreted by digital circuits as 16-bit words at 424 byte/s, or 32-bit words at 651 byte/s. This must be of concern to anyone with digital data equipment in the vicinity.

Nor should we treat all these health scares as urban myths. Sweden's MPT2 standard for magnetic interference sets stringent emission levels (2.5milliGaus) for ELF (300Hz to 3kHz) radiations — this is the area of the spectrum that gives researchers most cause for concern. According to NEC, some computer monitors show ELF levels 28 times the Swedish limit (generated by their horizontal deflection coils) — and these devices are reasonably well shielded. So this puts potential health dangers of ELF emission levels in perspective — surely you'd expect to find them in a TDMA radio transmitter.

You might argue that the Swedes are being over-cautious here, and that magnetic radiation is somehow more dangerous than radio waves, but you won't convince me. The fact is, that when we put a mobile handset against our head, we have no idea at all about its long-term effects. And high-energy pulsed R/F is likely to be more dangerous than constant low levels — this is the way chemists speed up some chemical reactions.

We don't need to start a scare campaign — but public fear is inevitable unless a few million dollars worth of money is poured into the problem quickly to prove, beyond any reasonable doubt, that this equipment is safe. Only in England are they doing that — it's a three-year project — and the dam of public ignorance won't hold for that long.

Preselection

Canberra Nominated as Ballot Testbed

Frank Blount, the blunt speaking CEO of Telecom, has been quoted as saying his job is to "lose market share gracefully." At the moment, on this measure he is not doing a very good job. Telecom has been repeatedly hauled before the Trade Practices Commission and given a telling off for publishing misleading advertisements comparing its long distance and international telephone rates with those of Optus. At the beginning of June it was forced to publish full page advertisements correcting the misleading information of previous advertisements.

Optus itself has not been squeaky clean in its own advertising, but at the time of writing had been spared the embarrassment of publishing retractions. And all this before the pre-ballot advertising blitz starts in earnest.

With so much misinformation being bandied, about the independent Austel-managed public information campaign on the preselection ballot will be crucial to giving the public some understanding of what the process is all about. No wonder Telecom and Optus could not agree on it.

Telecom objected to Optus' and Austel's proposals for the campaign and to the level of expenditure proposed ($5 million to be provided jointly by Telecom and Optus). Telecom pointed out that this cost would eventually have to be borne by the customer. This expression of concern is hard to swallow. Customers will eventually have to bear the cost of both carriers' advertising campaigns, and the estimates are they will spend over $100 million between them.

Austel was asked by Telecom to arbitrate on what should be the scope, thrust and cost of the ballot's public education campaign. Austel decided that the campaign should 'focus on creating awareness and understanding of the process' and decided that 'encouraging people to participate in the process is the responsibility of Optus and Telecom and not of (the) campaign.'

Austel completed its arbitration on May 31, removing the last hurdle preventing the ballot from starting. Even before this, Optus, which has the right to choose the timing of the ballot in each location, had announced its intention to hold the first ballot in Canberra on July 15.

This is something of a test market for the process. The Canberra telephone market is dominated by large Government departments — there are relatively few autonomous customers. The 1991 Canberra White Pages, for example, contains only 300 pages, compared to almost 3,000 for Sydney. By choosing Canberra, Optus will be able to fine tune its strategy and the consequences of any mistakes will not be too disastrous.

Austel said it had given careful consideration as to whether the public education campaign could be effective in the available time. The conclusion was yes, it could. It will be no small task. Austel has decided that the public education campaign must cover:

- The new Australian telecommunications environment;
- The reasons for and relevant characteristics, mechanics and implications of the preselection process;
- The transition process from Dial 1 Access to preselection, including the ballot;
- Services involved, services unaffected, fault reporting arrangements and billing;
- The options that are available to customers and the choices they will be invited to make;
- All opportunities to choose, including the use of the override codes;
- How customers might go about making their choices;
- The implications of not responding to the ballot;
- The treatment of silent lines; and
- Any privacy implications of preselection and the ballot.

Telecom did not think television advertising was necessary for the public education campaign. Austel did not agree: 'The campaign shall employ whatever media is considered appropriate to the various messages and target markets involved including television, radio, print media, direct mail, information booklets and 008 information lines . . . Public relations shall be used to promote community awareness and understanding and engender a positive community response.'

Some Helpful Advice

The public education campaign may have been the last point of disagreement between the carriers, but it was not the most important. Two issues in particular were fought over long and hard: the question of whether the 0015 fax service should be included in the ballot, and how calls initiated by Telecom's Easycall facility should be treated. The first was settled by Austel arbitration and the second by Ministerial intervention, and the entire process highlighted some fundamental issues for the regulatory structure.

Austel, in its arbitration, drew up some guidelines to be used in helping to resolve other preselection disputes. One of them said: 'Unless a carrier can exploit infrastructure ownership in some way, i.e. can gain a service cost advantage, a service differentiation advantage or a service access advantage over other carriers from investing in its own infrastructure, there may be little incentive in the longer term to do so.'

We've got to remember that cellular analogue dealers have occasionally reported that their customers claim headaches, eye-throbbing, and skin peeling off from hands, face and ears — and, until recently these claims were dismissed. Like most journalists, I've usually put these fears down as old wives tales, along with stories of aircraft crashing, car windows bobbing up and down — and in the case of GSM, BMW anti-lock brakes failing and air-bags blowing out.

But the research people looking into these out-of-band emissions don't take these stories as lightly as you or I; they remember that there were serious problems in the early days with car ignition systems and garage doors, so the anecdotes are too important to be ignored.

There's a warning and a disclaimer now being issued in many car manuals and with most of the new digital phones. They tell you not to use a digital handset with an 'integrated antenna' within the car. Tell that to the taxi-driver, whose passenger has just arrived from Germany and is making appointments from the back seat! You don't solve problems potentially this serious, this way.

At the high-frequency end of the scale there are also problems in Australia. Out-of-band emissions from GSM handhelds have been leaking into the adjacent cellular analogue band and disrupting MobileNet phones. We are the exceptional case here with GSM: we have our analogue handset receive-channels between 870 and 890MHz, and GSM's handset transmit-channels begin at 890MHz. In retrospect, this wasn't all that bright (although the Americans have gone one better, by putting their digital phones within the analogue band.)

What this lack of separation shows is that no one in Europe, Australia or the US twigged that GSM handset emissions were likely to be a problem, until it was too late.

In the interests of 'international roaming' and cheap handsets, we can't change that frequency selection now. So, according to TRL, we'll need to provide at least a MegaHertz or two of guard band between the phone systems, and therefore the five divisions of the GSM spectrum will now only possibly supply four. The 10:1 capacity gain originally promised for GSM has been whittled down now to 3:1 (claimed) and about 2:1 (in practice) — and now we've only got four-fifths of that. Not a good beginning!

Why the Rush?

We've got to look at what happened in two parts: Why did we rush into digital? And why did we choose GSM?

The government must take the blame for the digital push. Prime Minister Keating is a born-again digital convert. He thinks 'digital' is synonymous for 'smart' or 'lucky' — as in 'country' — so if it has the word digital

After issuing this arbitration Austel was asked to mediate on, among other things how Telecom's Easycall services should be treated in the ballot. Easycall allows the customer to program the local AXE exchange to dial numbers in response to abbreviated dialling codes and to divert calls to a specified number. Telecom argued that these calls should always go on its network, even if the customer had preselected Optus. Of course, Optus did not agree. The business involved is probably quite small, but it is an important principle for the reasons identified by Austel in its 0015 arbitration.

While Austel was in the process of mediating, Optus got the ear of the Minister and Austel put the mediation on hold until a definitive statement of Ministerial wishes arrived. The end result was that Easycall was included in preselection.

There has been much speculation as to what the Ministerial letter contained. Peter Leonard, a lawyer acting for Optus, quoted from it in a speech at ATUG'93. According to Austel Chairman, Robin Davey, the quote did not reflect the true gist of the letter. But it is instructive: 'There is no scope for market power at the local exchange level to be used to adversely affect competition in other parts of the network.'

Optus has announced no plans to install local exchanges, so will no doubt be very pleased with this decision. Optus has, however, decided to purchase DSC switches which will be installed in parallel with its Northern Telecom trunk switches to support intelligent network services. Discussing the sale with the author, DSC Vice President David Boyce said: "It is not possible to gain significant market share without a local exchange network." That, perhaps, is why a little Ministerial invention comes in useful.

Stuart Corner is the Editor of Exchange.

attached to the label, he can confidently make billion dollar technological decisions. We had to be the 'first country in the Southern Hemisphere with GSM,' just as we had to be the 'first in the Southern Hemisphere with digital Pay TV.'

And Telecom must take some of the blame also. It was originally given the B-band of AMPS frequencies, with the A-band reserved for a second mobile carrier. Back in 1987, it was widely anticipated that the licensing of a second mobile cellular operator would be the first stage of our move into duopoly telecommunications.

When MobileNet began to run into capacity problems in the larger cities, Telecom had two choices: a) reduce the cell-sizes, by building more base stations; or b) whack some A-band transmitters into the existing cell sites, and not tell anyone. The handsets

Continued on page 20

Malaysia

Quiet! Telstra is Bidding

Telstra's barely explained silence over its negotiations to join a consortium planning to launch Malaysia's second carrier appears to be more a result of Malaysia's peculiarly complex political economy than any innate publicity shyness on the Australian international carrier's part.

Telstra's fear is that if it is seen to be talking to the media at all, it will be seen to be breaking confidentiality clauses and thus jeopardising a Heads of Agreement it has apparently been negotiating with a Malaysian telecommunications licencee, Binariang Telecommunications, in a project estimated by Malaysian and Australian industry sources to be worth up to $500 million.

Binariang is controlled by colourful, high profile businessman, Tan Sri Ananda Krishnan, who, according to various media reports, has been granted licences to operate a GSM network, satellite services, and a fixed network in competition with the present monopoly carrier, Syrikiat Telekom Malaysia (STM).

Telstra has worked with STM to upgrade the management of its international telephone business. A Memorandum of Understanding was signed in 1990 to pursue joint venture opportunities in South East Asia. Under the MoU, Telstra has provided assistance in international accounting systems, a cost study of the international telephone service, computerisation of the foreign exchange ledger, and training on stream and product management.

But Telstra is reported to have rejected the idea of purchasing equity in STM, which is scheduled to be partially privatised.

The Telstra shutters appear to have slammed down hard in the wake of an *Australian Financial Review* article in late May, which reported that Telstra was set to win the contract with Binariang and effectively become the operator of Malaysia's second network licence. It was the up-beat, presumptuous slant of the story, rather than its substance, that apparently got Telstra's back up.

The report was quickly picked by the Sydney correspondent of Kuala Lumpur's *Star* newspaper, who quoted it in detail in an article published on May 22. Informed sources in Kuala Lumpur told *Australian Communications* the *Star* story had caused so much consternation that a Malaysian Government spokesman and other senior bureaucrats had publicly refuted it.

One Kuala Lumpur source said Telstra "will definitely be in on the new services. But the size of its stake or its influence will very probably not be as big as stated in the *Star* story. The Aussies won't have anything like the kind of role as second carrier network operator that the *Star's* story claims."

A Local Power Struggle

Another element affecting the Telstra/Binariang negotiations appears to be an intense conflict between Ananda Krishnan and rival Tajuddin Ramli, who controls Technology Resources Industries, a group involved in manufacturing and consumer electronics. Ramli is reported to badly want a stake in Malaysia's telecommunications future.

Estimated by the *Far East Economic Review* to be worth at least $160 million in late 1991, the 54-year-old Krishnan, known as the 'black jewel of Asia,' had business links with Australia through the failed Minsec mining operation in the early 1970s.

A former board member of the Malaysian central bank, Krishnan's diverse business interests include a satellite service beaming horse races into many Asian countries, oil and other resources, finance and manufacturing operations.

He is widely reported to have close links with both the Malaysian Prime Minister, Dr Mahathir and his family, as well as the Opposition. Krishnan is also closely associated with the national petroleum group, Petronas, which is overseen by the Prime Minister's office. Krishnan and Petronas are centrally involved in the multi-billion dollar Racecourse project, a commercial property development on 39 hectares owned by the Kuala Lumpur Turf Club in the heart of the Kuala Lumpur CBD.

"Both Binariang and TRI are vying for control of the second carrier. But where Tajuddin Ramli seems to be conducting a PR war, Krishnan is known to be a shrewd player who doesn't make any noise until he has a successful strategy in place," a Kuala Lumpur source says.

"Binariang has been awarded licences, but it's an open question whether TRI has been awarded anything that could amount to second carrier control. TRI, for example claims to have long distance and overseas capabilities, but that could be via some outmoded technology. "What we are witnessing is a publicity war fuelled by business and political rivalries," the source said. "There are four different licences in question, but on the technical and standards side, it's an eclectic hodge-podge at the moment.

"Also remember the satellite licence is worth hundreds of millions of dollars and the space agency regulators want to see secure and experienced operators backing any launch — especially after the Chinese disaster with Optus' B2 Hughes upgrade."

With no firm word on the Binariang/Telstra negotiations expected until the second half of the year, Telstra, meanwhile, is remaining mum, telling media representatives that anything is says can only be damaging to its Malaysian prospects.

Bernard Levy

Vietnam

Sanctions Hold Back a Would-be Tiger

Vietnam, once synonymous with war and political turmoil, is now flavour of the month with international business. The best hotels of Hanoi and Ho Chi Minh City are full of foreign executives swapping rumours about the impending end of the US embargo.

Vietnam, the reckoning goes, will be the next Asian tiger. Despite being one of the poorest countries in East Asia, with per capita income of less than $US200, it boasts a huge local market (60 million), eight per cent annual growth, 80% literacy and a Confucian work ethic.

Amid all this, the best-established Australian company is Telstra, a position confirmed by the visit of Prime Minister Vo Van Kiet to Telecom's Sydney network management centre during his recent official tour.

Since the then OTC International built Intelsat earth stations in Ho Chi Minh City and Hanoi in 1986-7, Vietnam has become Telstra's Asian beach-head as it gropes its way towards becoming 'the leading regional carrier by 2000.'

Telstra is a quarter way through a ten-year Business Cooperation Contract with the Directorate-General of Post and Telecommunications (DGPT). Under this, Telstra must spend $US66 million on the national network and another $US5 million on training. In return, Telstra earns a share of international call revenues. Despite the increase in call volume since 1988 from four million to 40 million minutes, the company had spent $US40 million to the end of 1991 and is still in payback period, according to Indochina General Manager John Malcolm.

A joint Telstra-DGPT Advisory Management Committee determines priorities for network development and manages the international business. From Telstra's point of view, network expenditure must be *viable* — in other words, likely to increase international traffic. For example, in 1991, Telstra bought a new switch for the main northern port city of Haiphong.

The good relationship with the DGPT provides a platform upon which Telstra can build its business. Activities and plans at the moment include:

- The 565Mbps Thailand-Vietnam-Hong Kong (TVH) optical fibre cable, construction of which will begin later this year. Telstra is the co-signatory with the DGPT for the Vietnam leg.
- Compilation of a network plan for Ho Chi Minh City and neighbouring areas, the source of more than half of GDP.
- A new headquarters in Hanoi for the DGPT and its foreign suppliers and clients is under consideration. Vietnam Operations Manager, Maurice Sparkman, said a 10 storey intelligent building was envisaged. The current offices are cramped, with uncertain power supplies and limited access to the local loop.

Like everyone else, Telstra also must endure poor quality local loop access to the national network. The aging local loop is blamed for the slow progress of Telstra's packet switched service, Data Access, launched in late 1992.

Despite Telstra's successes so far, Vietnam is a difficult market in many ways. Only one telephone is available for every 400 people, although the DGPT target is a four-fold increase by 1995 to one for every 100, and social traffic still exceeds business traffic in the international network. The US sanctions prevent Vietnam from raising loans from international development funds such as the World Bank or Asian Development Bank. The rumours regarding the lifting of sanctions have been current for several years. No-one really knows, but certainly every major international telco — including AT&T — has beat a path to the DGPT door.

Restrictions from the Western alliance technology body, Cocom, are also enforced. As a result, the national optical fibre backbone is restricted to 34Mbps, rather than, say, 565Mbps, and Telstra has not been able to deploy the caesium clock required to synchronise timing in the network.

The network itself is a welter of different components. The process of digitalisation over the last six years has seen the introduction of switches from most of the major suppliers — Ericsson, Alcatel, Siemens, NEC, and Fujitsu among others. As the network evolves, integration of the various parts will become a minor nightmare, Sparkman says.

The TA(I) Deal That Wasn't

Finally, bureaucratic uncertainty and rivalry rank high on the list of frustrations for foreign businessmen. A classic case involves Telecom Australia (International). The overseas arm of the pre-merger carrier announced in November 1991 it had won the right to build and operate a cellular network.

True, TA(I) did get the go-ahead from the DGPT, but at the same time the Ho Chi Minh City PT office awarded a similar contract to arch rival Singapore Telecom. The size of the Ho Chi Minh market rendered the TA(I) contract meaningless.

Yet, Malcolm explains, Telstra is not out of it yet. The DGPT is is keen to develop a digital mobile service, and is still tossing-up between pushing ahead with AMPS under Singapore Telecom, or throwing open tenders for a national digital network.

Robert Clark

Multimedia

US West Buys into Time-Warner

A major US telephone company and the entertainment conglomerate Time-Warner have joined together to try and provide "a one stop shopping source for local cable (TV) and telephone service." Those are the words of US West Chairman, Richard McCormick, whose company, US West, is acquiring 25.5% of Time-Warner. That 25.5% will cost US West $US2.5 billion. In addition, US West will spend $US5 billion to upgrade Time-Warner's cable network to optical fibre over the next five years.

Industry analysts say this is a unique alliance for many reasons. Time-Warner owns an unrivalled library of movies and other entertainment that US West will be allowed to use on its networks.

Telecommunications analyst with TeleChoice in New Jersey, Daniel Briere, said there are three major areas emerging for consumers to receive information: cable, telephone and wireless/satellite. This alliance will open up competition for new services, which means an abundance of choices and lower costs for consumers.

"Time Warner is *the* programming entity with lots of entertainment choices to offer and it has teamed with the right company who can push those choices," Briere said. "Over the next five years there will be a true explosion, and consumers will start seeing the benefits and the options."

US West said that the transaction has been structured to conform with the Modified Final Judgment regulating activities of the Regional Bell Operating Companies (RBOCs), of which it is one. Separate entities will be created in which US West has no financial interest, to enable the company to own certain of Time-Warner's facilities and other assets that it may not be able to hold under the current regulations.

US West will also request a waiver from Judge Harold Greene, overseer of the decree that split AT&T and the RBOCs, so the assets can be integrated with Time-Warner.

As well as the usual catalogue of future services (like home shopping) trotted out on occasions such as these, US West says that its users may eventually be able to access each book in the Library of Congress.

One thing is certain; with US West and Time Warner hooked up, there won't be a need for video stores because users will be able to download movies and have the ability to pause, rewind and other such functions, all with the touch of the remote control button, Briere said.

Ann Steffora

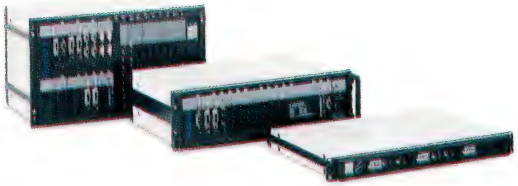

Stock markets go up and down, but your network can't. That's why LANNET designed MultiNet™ - intelligent hubs with

THE MARKET IS UP. UH-OH! *Our network is down*

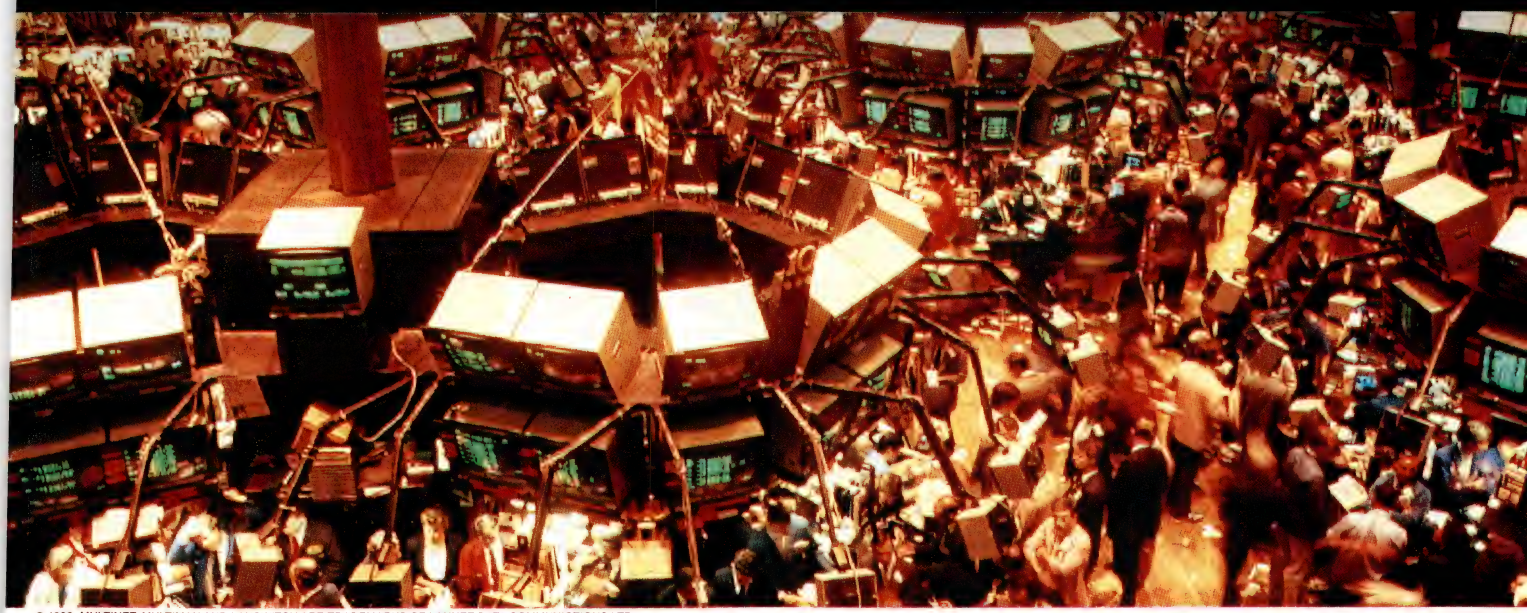

blue-chip reliability. Robust design includes port, segment, and power redundancy with no single point of failure.

Failsafe - and flexible. Monitored modules **FAULT TOLERANCE** *bring multiple-network, multisegment Ethernet, Token Ring, FDDI, LocalTalk, and RS-232 into one hub, with integral bridging and routing. The MultiMan™ SNMP management system provides powerful centralized analysis tools.*

Protect your networking futures. MultiNet hubs incorporate LANNET's LANswitch,™ an ATM-ready 1.28 Gbps backplane, delivering bandwidth for full-motion video and workgrouping over today's 10BaseT networks.

If you're looking for a secure invesment with high dividends, it pays to call LANNET. We're one market that won't let you down.

Data Comm MAGAZINE HOT PRODUCTS 1993 AWARD
LANNET LET-36 WINS!

TOREN Member of the Ipex Group
Melbourne: Tel / Fax: 61 3 242 5055
Sydney: Tel: 03 319 2755 Fax: 02 699 8004

Dataplex
Melbourne: Tel: 03 210 3333 Fax: 03 210 3399
Sydney: Tel: 03 3493911 Fax: 02 349 7747

LANNET

THE FAILSAFE FUTURE OF NETWORKING

INT'L. HEADQUARTERS: TEL 972-3-6458458/ FAX 972-3-5447146 U.S. HEADQUARTERS: TEL (714) 891-5580/ FAX (714) 891-7788 FRANCE: TEL 33 1 69 85 24 11/ FAX 33 1 60 19 29 54

Asia Telecom 93

Where Were the Australians?

Depending on your perspective, the Asian telecommunications market is either a goldmine of opportunity (vast populations with almost non-existent telecommunications infrastructure) or an insurmountable problem (nations struggling to develop modern economies but with almost non-existent telecommunications infrastructures and no resources with which to develop any).

Attention was focused sharply on both facets of the market by several events associated with the International Telecommunication Union's Asia Telecom 93 conference and exhibition held in Singapore at the end of May.

For Australia, with its eye on Asia and with telecommunications seen as one of its major export earners, the event should have been a great opportunity to put on a high national profile. It was not exploited. There were 14 national pavilions, Australia's was not one of them and there were only a few Australian organisations with stands of their own, including Telstra, Olex Cables and *Australian Communications*. The New Zealanders put Australia to shame with a good sized national pavilion and eight participating companies, including one Australian-owned — Exicom International.

With a recent name change to explain, the Telstra stand in particular was something of a disappointment, and failed to match the scale of CEO Frank Blount's Asian vision in either size or impact.

The ITU, just before Asia Telecom 93, held its first Asian Regional Telecommunications Development Conference. Two developing nations that discussed their telecommunications plans were India and Indonesia. Both saw allowing private investment in their networks as an avenue to funding much needed expansion.

India's Communications Minister, Sukh Ram, told a press conference that India would consider allowing private investment in its public network in order to help it finance its plan to install at least one payphone in each of 225,000 villages by 1995. Jonathan Parapak, Secretary of Indonesia's Department of Tourism, Post and Telecommunications told the conference that his Government would introduce a telecoms duopoly in a bid to get private capital to fund network expansion and upgrading.

The Development Conference produced a long list of recommendations and resolutions which might well constitute a wish list for those companies with an eye on market opportunities in the region. It recommended

- Separation of regulatory and operating functions;
- Adoption of an investment base as wide as possible;
- Introduction of intelligent networks in each country of the region;
- Access to telecommunications for all, at least to basic telecommunications services, in a competitive environment;
- Local manufacturing of telecommunications equipment and components whenever appropriate;
- Harmonisation of networks to promote global inter-operability and to strengthen integration;
- Use of telecommunications for protecting the global environment

It recommended countries develop a timetable for:

- Providing at least one public phone in every village including national/international dialling facility;
- Providing public pay phones extensively in urban areas;
- Setting up telecommunication service centres for fax, telex and other value-added services in every business/commercial community in areas where economic activities would be enhanced; and
- Increasing radio an TV coverage to rural areas.

How, when or indeed whether these ambitious goals will be realised is entirely another question. As ITU Secretary General Pekka Tarjanne said in his closing address: "The Conference will not be worth much unless we ensure that the follow-up mechanisms are efficient and able to produce real results — concrete telecommunications development for the people of this region."

Stuart Corner is the Editor of Exchange.

Exports

Alcatel TCC's South American Coup

Small wonder senior management at the Alcatel Tasman Cable Company (TCC) seem to exude a kind of serene inner glow. Any company poised to take a solid slice of an estimated $10 billion-plus market before the end of the century would have good cause to count itself lucky.

When the Sydney-based company recently announced it had been awarded a $33 million contract to manufacture 1,200kms of submarine optical fibre cable to link Brazil, Uruguay and Argentina under the $76 million UNISUR project, it highlighted its export earning potential. While much of the profit from its annual $200 million turnover is exported (it is owned one-third by Alcatel Australia, with equal shares held by Alcatel Cable and Alcatel CIT in France), the company points out that it employs 340 staff at its two Sydney plants and another in New Zealand, buys most of its materials in Australia and spends large sums with local subcontractors and secondary industries.

The 1,200kms of cable for the UNISUR project will be produced at Alcatel TCC's $120 million facility at Port Botany. The company will also have overall responsibility for integrating the South American system, which also incorporates repeaters from consortium partner AT&T, and an additional 600kms of cable from Pirelli in Italy.

But the $33 million contract is literally the tip of the iceberg for Alcatel TCC. General Manager Michael Kerr estimates that demand for submarine cable in the Asia-Pacific region, which is projected to account for half of the world telecommunications market before the end of the decade, would be "easily $10 billion by 1998.'

Alcatel TCC has the competitive advantage of being close to the Asia-Pacific market in a four-way contest with AT&T, Japanese consortia and STC in the UK, to supply turnkey cable systems, Kerr says. "We also have the advantage of being the world's newest submarine cable producer, and we were able to build our Port Botany facility up from the ground, rather than converting its production from analogue coaxial cable, as the Alcatel Cable plant in France had to do."

In the UNISUR project, the cable will originate in Brazil, travel 1,500km up the South American east coast, then split via a branching unit to Uruguay and Argentina. When completed in August next year, the cable will connect with the AMERICAS-1 and COLUMBUS II cables linking South America to the global networks.

Built with a 25-year reliability factor at the bottom of the ocean, the Alcatel TCC cable works out at $20,000 per kilometre. The company's involvement with UNISUR means it will have supplied 17,000kms of submarine cable, including the Tasman 2 link between Australia and New Zealand, the PacRimEast and PacRimWest projects, as well as the SEA-ME-WE2 link between Asia, the Middle East and western Europe.

Michael Kerr sees enormous business potential among developing nations looking to link into the world's new generation, high traffic backbones. "The UNISUR project, for example, is the first step by South American nations to link into the global digital optical fibre networks, and this trend will continue around the world," Kerr says. "The big telecommunications markets established new generation links and now the smaller nations realise that to capitalise on the general growth in telecommunications, they must connect to the upgraded network. We're perfectly placed to help them do it."

Bernard Levy

Open access architecture. For an ever expanding universe.

In the rush to embrace the benefits of faster communications platforms, you may also have noticed a drawback. New emerging protocols often require a new analyzer. And that can be costly.

So Tekelec developed a solution. Open Diagnostic Access. An architecture that shares a common platform and provides a modular approach to expanding protocols. This means your diagnostic and monitoring needs can be met for now and the future with any one of a family of analyzers, depending upon your primary needs.

That's because Tekelec is the only company with products that analyze recently developed protocol standards on all key technologies: Frame Relay, FDDI, ISDN, X.25, Ethernet, Token Ring and SMDS (both DS-1 and DS-3).

Now you can monitor multiple LAN and WAN networks simultaneously. The addition of a new technology means a significantly lower capital outlay. Commonality of products ensures lower training time and costs. And the X Window based architecture is naturally suited to distributed diagnostics.

Open Diagnostic Access. It's the kind of thinking you'd expect from Tekelec. After all, Tekelec has been the innovator in protocol analysis since 1979.

Call TEKELEC today.

TEKELEC

Tekelec Australia Pty Limited
Suite 14, West Court, Bankstown City Capital Centre
41-45 Rickard Road, Bankstown NSW 2200 Phone: (02) 790 2200 Fax: 790 0848

Telecommuting

Telecentres Boost Distance Learning

Telecom Australia will have helped create a new telecommunications market the size of a medium-to-large Australian city by 1996, if the growth of regional telecentres, or 'telecottages,' continues at its present pace.

The Department of Primary Industry and Energy (DPIE), the major source of government funding for the telecentre scheme, estimates the number of clients using the centres is already around 50,000. But by 1996, when the last of the $2.8 million allocated in the August 1992 Budget has been spent on establishing some 40 centres, the number of users would be in the hundreds of thousands, according to DPIE's Telecentres Program Administrator, Jim Graham.

"We will be gathering hard statistics as the scheme develops nationwide," Graham says. "It is likely, in the 'clever country' context, that more money will be made available in this year's Budget."

The telecentres provide users with telephones, faxes, computing and audio-conferencing and distance learning facilities designed to increase computer and telecommunications skills and job opportunities in rural areas. The centres also work closely with the DPIE's Countrylink service, which provides rural residents with information on such subjects as the weather and land and soil conditions via a 008 telephone number.

The number of telecentres has increased in less than a year from the original three — one at Cygnet in Tasmania; one at Byron Bay in northern NSW; and another at Walcha in north western NSW — to nine, with proposals for another five now being assessed by the DPIE for funding this financial year.

In the latest round of funding, a total of $430,000 has been allocated to two of the existing centres (Byron Bay [$36,000] and Walcha [$40,500]) and for the establishment of six new centres — at Blackall in Queensland ($61,000); Merredin ($52,000) and Wongan Hills ($50,000) in Western Australia; Wangaratta ($116,000) and Maffra ($30,000) in Victoria; and Wudinna in South Australia ($51,000).

The Wangaratta project will be the first regional telecentre network, with several small shopfront operations in local towns linked to the main centre.

While the government funding goes towards establishment costs such as premises, hardware and the salary of a coordinator (usually around $20,000), the centres are expected to be financially self-supporting within two years, selling their facilities and services to the local community at moderate rates. In Byron Bay's case, for example, the centre offers computer courses to the unemployed through the CES and local Skillshare Centres, as well as printing and distributing a newsletter for the local Chamber of Commerce. It is also seeking expressions of interest from local schools.

Typically, there are two models for the centres — those formed under incentive schemes run by the various by state governments, and others from ad hoc community groups of activists and social networkers.

While the DPIE has overall responsibility for the scheme, with support from the Department of Employment Education and Training (DEET), Telecom Australia can take much of the credit for initiating it.

Telecom's Manager, Rural and Remote, Consumer National Office, Tom Cass, says Telecom has deservedly reaped the benefits of adopting a participatory approach and looking at the rural market from a 'consumer needs' point of view, rather than a technological approach based on selling hardware.

"The Strategic Planning section of Telecom's Country Division went to the rural areas and asked people what they needed in terms of new communications. We also worked with the DPIE and researchers at the University of New England, who had identified low computer skills as a problem in rural areas generally, not just in the agricultural sector," Cass says.

"It took a while for the original Scandinavian concept to transpose itself into an Australian cultural context, but it has gained its own momentum now and hundreds of people are gaining skills which will allow them to find work." Cass says that ideally, the telecentres would all eventually have a dual distance-learning/distance-working role, promoting the concept of 'telecommuting' to work, providing communications services for local schools and tertiary institutions and, to some extent, replacing correspondance courses.

Cass says the very latest trends in the development of telecentres will be aired at an international conference, Telecottage 93, to be held at the Gold Coast from November 29-December 1, and featuring speakers from Sweden, the UK, the US, Brazil, Hong Kong, New Zealand and other countries.

Bernard Levy

GSM from page 15

could handle both. Guess which approach Telecom chose?

The government caved in and retrospectively licensed Telecom to use both A- and B-bands to save face — and then later they had to quieten an angry Optus which didn't have a band of its own. So it hurriedly decided on the 'first in the Southern Hemisphere' strategy, and ordered that analogue cellular should be wound down by the year 2000. That served it up to Telecom, but tough luck for the users!

By the time Vodafone came along this year, it was also apparent that Telecom was thinking seriously about double-dipping into digital by slipping some Digital-AMPS and/or CDMA transmitters into the MobileNet system. Both these systems reuse analogue frequencies.

As part of Vodafone's licence agreement, the government has now stipulated that digital mobile telephony cannot be used in the analogue band. Effectively, this over-reaction has now cut off the most promising of the new digital technologies (CDMA) — unless there's a reversal of policy, and probably a long legal bun-fight.

The choice between GSM and D-AMPS (CDMA wasn't ready at that time) came about for philosophical, rather than technical reasons. According informed sources at Austel, Alex Arena, who wrote the recommendation, was determined not to restrict digital mobile to a duopoly. D-AMPS (which has its own TDMA-related problems) allows bandwidth for two operators, while GSM goes the whole hog with five (now four!).

I guess we should also ask why the GSM decision wasn't later reversed when problems arose. It has been apparent for at least two years that GSM was in trouble. Consider the evidence.

There were the constant rounds of promised roll-out dates and subsequent delays; the most unlikely claims were made that delays resulted from a lack of type-approval equipment; then there was the fiasco over the patent-pool which never eventuated; the play-down of claimed capacity increases; the lack of technical papers showing the success of trials; roll-outs of supposed networks that never had any real commercial users; the splitting of GSM's features into Phase-1 and Phase-2; the projected delays of Phase-2 until 1995-96; headlined stories in the technical journals saying 'Hopes For GSM Fade' and stating 'Quite simply: GSM has been recklessly over-engineered. And Europe is now paying the price;' the revelation that voice-quality tests showed GSM introducing 7QDUs of distortion; leading up to the A5 encryption cock-up.

Alvin Toffler pointed out that individuals were having trouble handling the pace of technological change. It's now becoming obvious that our political and cultural institutions are having even greater problems.

But when, worldwide, technocrats, scientists and technologists aren't capable of keeping control and making rational judgements on these complex developments, then we've got serious problems. This is not to suggest that there's anything wrong with the people, it is the institutions and the structures of decision-making that are failing.

Stewart Fist

Now available in red.

Now you can get the leading integrated solutions in the most popular networking color. Novell® red. Because SynOptics products are now tightly integrated into the Novell NetWare® environment, including the NetWare Management System™ (NMS).

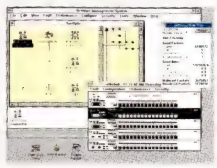

Optivity™ for NMS™ means easier management of Novell® NetWare® networks.

With network management solutions like our Optivity™ for NMS,™ you get an expanded view of your network, plus access to other NMS services. And SynOptics' Lattis EZ-View™ software, ideal for smaller networks, has Novell IPX support.

SynOptics intelligent hub products—whether Ethernet or Token Ring—are easily integrated into Novell NetWare environments. And because they're bilingual, they can be managed through SNMP over IP or IPX. So now you can build a flexible network that will grow as you do.

yes It runs with NetWare

Plus, our membership in the Novell Technical Support Alliance, along with our LattisSupport™ service products, provides comprehensive service from planning to installation and maintenance.

By supplying the solutions you need, SynOptics is helping you build the foundation for a strong network fabric—the underlying structure that connects and manages your network. It provides high-speed network communications. And it's flexible enough to grow as your network does, without sacrificing your existing investment.

To learn more about SynOptics' Optivity for NMS solution, call Australia 008 81 7070 or New Zealand 0800 44 7070. And see just how good we look in red.

SynOptics

The Network Fabric of Computing

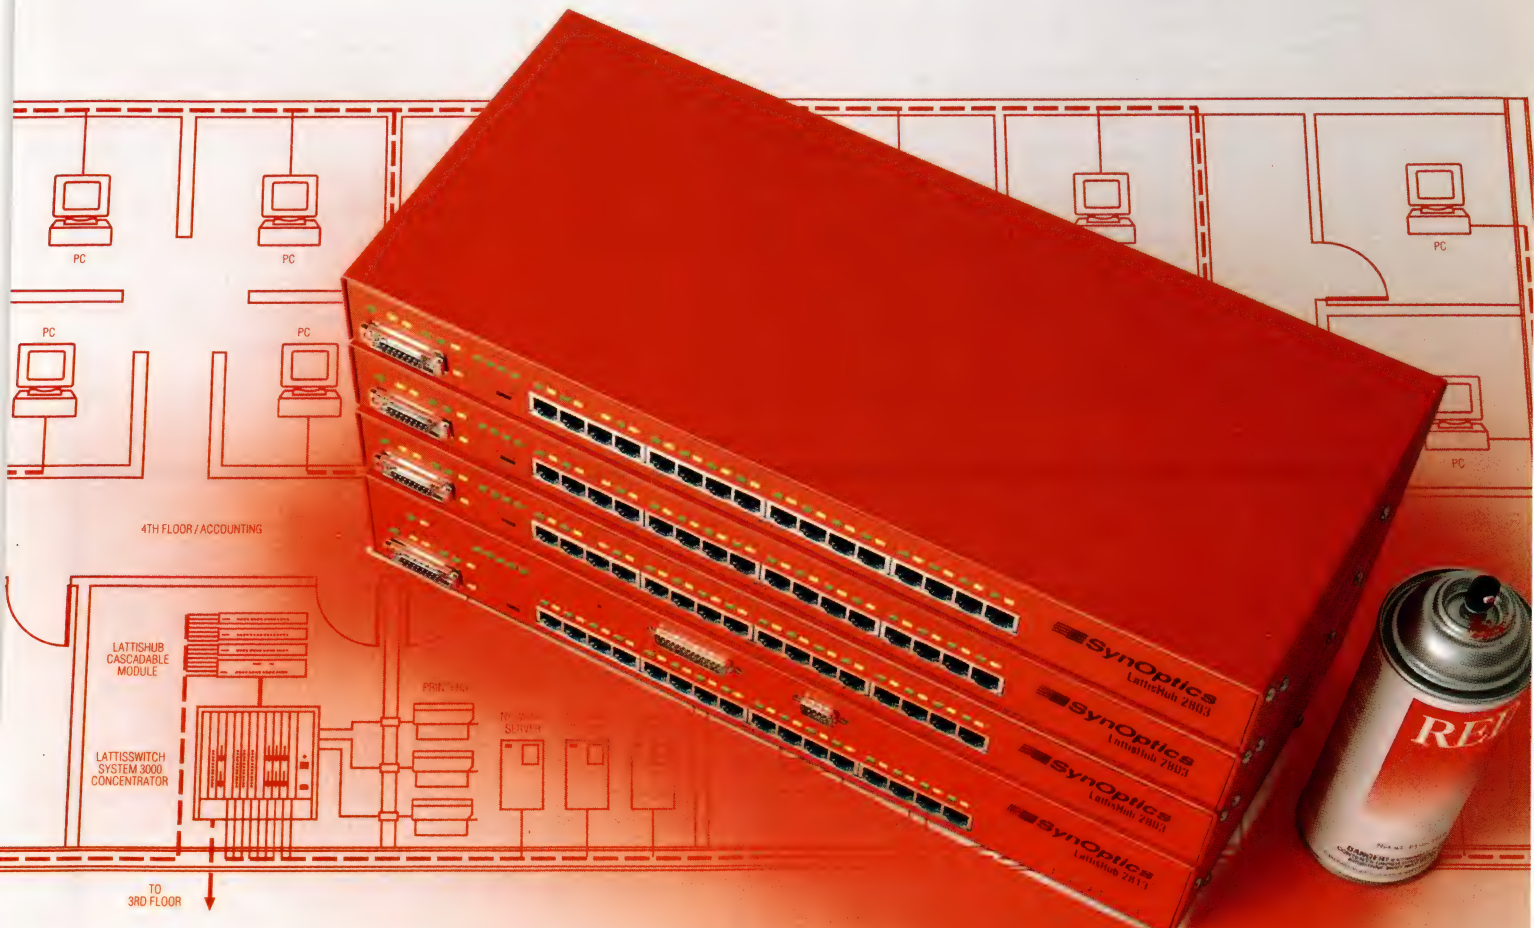

Global Communications

Carriers Jockey For Position in the Global Goldrush

The winds of change generated by the deregulation of telecom monopolies around the world are starting to blow strongly in favour of international network managers. After years of enduring abysmal service and oppressive prices, managers can look forward to much-improved conditions. And they can thank the global carriers building worldwide backbone networks that aim to deliver to corporate customers much better services at lower prices.

The charge toward global network services is being led by AT&T and BT, the two carriers most profoundly affected by telecom deregulation up to now. Faced with intensifying competition at home, both have devised aggressive plans to build global overlay backbone networks to offer the world's largest organisations alternatives to building and operating private international networks.

BT has just announced a deal with MCI Communications Corporation which will see the two companies form a new joint venture operation to provide global seamless value added voice and data services to multinational customers. BT and MCI will both market new services and products developed by the joint venture, and transport traffic from these services via a global intelligent network to be established in the near future.

Meanwhile, AT&T has announced its own new alliance, to be known as WorldPartners. Along with AT&T, other founding members are Japan's Kokusai Denshin Denwa and Singapore Telecom. Telstra, Korea Telecom and Canada's Unitel are set to join as associate members in the near future, with other carriers expected to follow. WorldPartners will market a new portfolio of network services called WorldSource, which will include frame relay, managed private lines for voice/data, and a switched virtual private network. The WorldSource service is being launched in the US and the Asia Pacific region this year, and in Europe starting next year.

For both AT&T and BT, global network services are part of an overall strategy to become major forces in what is expected to be an incredibly lucrative market: the international outsourcing business. The thinking is that as carrier competition heats up on the international level, profit margins for providing raw bandwidth will shrink, with more spoils going to companies that can provide a broader range of services. AT&T has been busy buying a collection of value-added network (VAN) providers and other specialty vendors with an eye toward developing total systems integration and facilities management offerings to entice global users. BT also has dipped into the VAN pool and is putting its integration chips on Syncordia, its outsourcing subsidiary.

Although AT&T and BT are first off the mark with global service offerings, they will not have the market to themselves. Other PTTs, some of which see AT&T and BT as trying to steal their best customers, are scrambling to form cross-border alliances in an effort to become global carriers themselves. VAN providers are expected to join the fray, although to a more limited extent for the time being. VANs lack the network resources to set up global backbones that can compete on price with those from carriers that own the cables, but their value-added expertise may come in handy as demand grows for international outsourcing services. With deregulation, the way is opening for VANs in Europe to carry voice as well as data traffic. Existing systems integrators and specialist companies also will be competing for outsourcing dollars.

Not all carriers are convinced that the aggressive outsourcing plans of AT&T and BT will pay off. Britain's Cable & Wireless, which now has a networking presence in around 40 countries, contends that no one company can hope to fulfill a total systems integration role like that envisioned by its rivals. C&W says it will leave areas like applications to more expert hands.

Managers of international networks now stand to gain some immediate benefits from all this carrier jockeying. As the new global carriers start competing with one another, charges should fall fast. The fully meshed, centrally managed backbone networks being installed by these carriers should deliver much better service than existing networks, which typically are patchworks of dissimilar national networks operated by various different PTTs. Competition should slash provisioning and repair times, and the ability of global backbones to re-route traffic automatically around failures should bring much greater reliability.

For carriers, the move toward global services makes sense for two reasons: Most major corporations now have a global reach, and the field for providing international networking services is wide open, at least for now. International VANs have offered services over centrally managed backbone networks for some time, but under previous restrictions they were not allowed to offer users clear channels or to carry voice traffic. PTTs and national carriers had grown used

OpenGate™ is king of all networks.

Modular Multiprotocol Router

In today's multimedia enterprise-wide network you need OpenGate™. The only multiprotocoal router that delivers parallel RISC technology and total fault tolerance. There's an OpenGate™ platform that's right for you, from a small office router to a 12 slot multiport backbone router. You get redundancy and hot-swapping for all critical hardware, as well as field proven traffic control features that improve network operation. With OpenGate™, it is clear why RND continues to be one of the world's leading pioneers in internetworking.

To find out more about OpenGate™ call RND today. You'll be surprised how affordable is the king of all networks.

OpenGate™

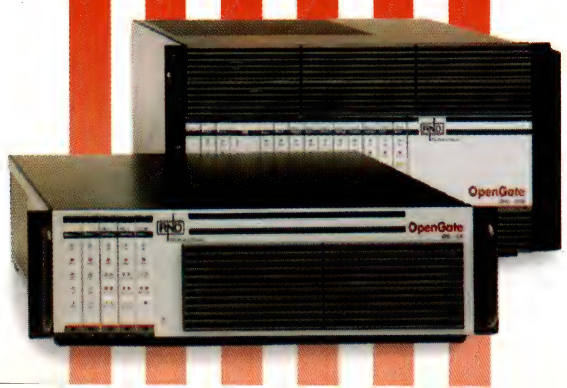

RND For Networks With Ambition

RAD Network Devices

Western Region
7711 Center Avenue, Suite 270
Huntington Beach, CA 92647
Tel: (714) 891-1446
Fax: (714) 891-1764

Eastern Region
151 West Passaic Street,
Rochelle Park, NJ 07662
Tel: (201) 587-0404
Fax: (201) 587-0334

TOREN Member of the Ipex Group

TOREN Computer Communications Pty,Ltd.
(Member of the Ipex Group)

Melbourne
Tel: (03) 242 5050
Fax: (03) 242 5055
Sydney
Tel: (02) 319 2755
Fax: (02) 699 2093
A.C.T
Tel: (06) 253 2665
Fax: (06) 253 2669

Dataplex

Dataplex Pty, Ltd.

Melbourne
Tel: (03) 210 3333
Fax: (03) 210 3399
Sydney
Tel: (02) 349 3911
Fax: (02) 349 7747
A.C.T
Tel: (06) 285 3657
Fax: (06) 285 3696

to charging excessive rates for international lines and using those profits to subsidise domestic traffic. With deregulation, they no longer have this luxury.

In Europe, the impetus for deregulation has been implementation of the so-called services directive from the Commission of the European Communities. The directive called for incorporation of relaxed restrictions in national laws by January 1, 1993.

Other countries around the world have made similar changes to their regulations in the past year or so. One of the more critical developments is now playing out in the US. BT announced earlier this year that it has applied to the Federal Communications Commission for a common carrier licence to operate an international virtual network service. Under existing regulations, operators of global backbones are allowed to carry voice traffic on point-to-point links, but they aren't permitted to switch that traffic at foreign locations. In BT's case, it has to hand off voice traffic destined for the US to an American carrier. Since making its licence application, BT announced its strategic alliance with MCI, and at the time of going to press it was still unclear what effect, if any, this might have on the FCC's decision. If the FCC grants BT a licence to switch voice traffic in the US — an outcome that is uncertain at present — it could open the way for national carriers to compete directly with one another for voice traffic.

Under the relaxed rules already on the books, VANs are allowed to offer clear channels over their international networks, but price competition from carriers that own the network pipes will make bandwidth reselling only marginally profitable at best. Instead, it's more likely VANs will stick to offering value-added services, such as packet switching and messaging. Because of this, it is the national carriers which are expected to dominate the emerging global carrier business. But competition among these carriers is sure to squeeze profit margins for selling raw bandwidth. US carriers already are 'barely breaking even' on the half circuits they supply on international routes, according to industry sources.

Carriers also expect corporate users to abandon private networks in favour of services that are tailored for their use. Rather than sit back and allow systems integrators and VANs to corner this market, carriers are angling to become outsourcers themselves.

Getting Ready

AT&T and BT have been particularly aggressive in laying the groundwork for global carrier services. Both have spent heavily to build global backbones and to get up to speed on outsourcing by purchasing companies specialising in systems integration.

AT&T's outsourcing plan includes its purchase of VAN provider Istel. The company, now known as AT&T Istel, is building a European backbone based on multiple 2Mbps circuits (see table on page 27). AT&T has also purchased facilities management companies in France and Germany. Its acquisition of NCR gives it access to potential customers in corporate data processing centres as well as telecom departments.

For now, the two carriers have few customers to show for their money, but are quick to point out that they are still in the early stages of building their global services. However, some evidence suggests that progress so far has been disappointing. For instance, AT&T Istel recently dismissed 375 staff in a belt-tightening move. "We geared up for really terrific growth last year," said an AT&T Istel official. That growth has yet to materialise.

BT must also be disappointed by the number of customers using its global backbone network, which is being built by its Syncordia subsidiary. At press time, Syncordia had only four customers, but said a half-dozen other contracts were in the pipeline.

Syncordia won't divulge any details about its backbone, but the network supports BT's Managed Links clear-channel service, which gives some clues about its scope and coverage. Bandwidths of up to 384Kbps are offered on the Syncordia network.

Along with pumping development funds into Syncordia, BT plans to expand Tymnet, the US VAN operator it acquired in 1989. Tymnet's products are now sold under the banner of Global Network Services (GNS). BT also is counting heavily on its aforementioned FCC licence request to boost its global carrier efforts. Its planned international virtual network service, known as the Cyclone project, takes the concept of a global backbone a step further by adding switches that will enable BT to offer virtual private networks for voice as well as data.

A Smaller Bite

Some competitors say BT's moves amount to a technology crapshoot. According to these detractors, BT is using high profits from its status as the dominant common carrier in the UK to engage in some speculative network building in the hope of attracting customers.

Would-be global carrier C&W contends that AT&T and BT are biting off more than they can chew by trying to offer facilities management as well as network outsourcing. In some ways, C&W is closer to becoming a global carrier than either AT&T or BT. It has an ownership stake in more than 40 carriers around the world, including ones in all of the countries acting as major hubs

AT&T SYSTIMAX Premises Distribution System.

It connects everything you can imagine. Even things you can't.

Imagine a wiring system that connects all the different systems that exist in today's buildings. Now imagine it can also connect systems that don't even exist. Yet. That's AT&T SYSTIMAX® PDS. It allows you to integrate equipment from vendors as diverse as DEC, IBM and HP. It even allows technologies like ISDN and FDDI to work together. Seamlessly. Transparently.

AT&T SYSTIMAX PDS is an advanced, modular, building wiring system with an open architecture that combines both copper and optical fibre with a 5/15 year assurance programme. With SYSTIMAX PDS, you're ready for everything you can imagine. And a lot of things you can't. No other system offers you greater flexibility. SYSTIMAX PDS. Call AT&T Network Systems at 008 804 738.

AT&T
Network Systems

ATT3008.FPC.R

for international private networks. The company already has built a Global Digital Highway comprising fibre networks in each of these countries, linked by undersea fibre cables. It has extended this network across Europe, linking 21 cities, and is offering a range of Customised Network Services, including clear channels and voice and data networks, all with centralised management.

Carrier Collaborations

Although it's too early to tell whether AT&T and BT will succeed in their grand outsourcing efforts, competing carriers are reluctant to let them get too far ahead. The pressure to keep within reach is particularly acute for AT&T's domestic competitors, Sprint and MCI. AT&T's rivals still have to focus much of their attention on building up their domestic businesses, which means they don't have such deep pockets for international expansion. Sprint also can't risk offending PTTs by aggressively competing with them, because a lot of PTTs are buying packet-switching gear from Sprint. But if MCI and Sprint fall too far behind in the global carrier business, they may risk losing domestic customers who want international services.

This dilemma has led both carriers to follow a strategy of collaborating with PTTs rather than competing with them. MCI's recently announced joint venture with BT is a case in point, while Sprint is developing its own international backbone and forming alliances with PTTs and national VANs that use Sprint packet equipment.

MCI, meanwhile, is also working with PTTs to build a global backbone while leaning on the expertise of an international VAN for outsourcing services. MCI owns 25% of Infonet, whose other shareholders are PTTs. MCI founded the Financial Network Association, a multicarrier alliance that is building a global backbone.

PTTs around the world are in near panic at the prospect of seeing AT&T and BT raid their corporate customer bases. Some recognise that they're too small to defend themselves alone and are going the collaboration route in the hope of finding a winning formula. One of the furthest advanced PTT collaborations is Unisource N.V., which was set up by the Netherlands and Swedish PTTs, and may soon include the Swiss PTT as an equal shareholder.

Unisource already has more customers than BT's Syncordia, although most of them were inherited from its shareholders' data service divisions, which have been converted into Unisource subsidiaries. Unisource is building a European backbone and is already offering services over it, including a clear-channel service called Unistream.

But plans to deliver packet-switching, frame relay, and messaging services appear to be slipping, possibly because of equipment incompatibilities in the Dutch and Swedish national networks.

FNA Under Way

Meanwhile, the MCI-led FNA collaboration has just completed its global backbone, and has begun to offer services on a trial basis. Several PTTs have applied to join the original 11 founding members, including the Swiss PTT, Telecom Denmark, Telecom Ireland and Telecom NZ. FNA services are marketed by members, whose contributions to the organisation's budget include network and staff resources as well as money.

As its name implies, FNA is targeted at banks, brokerage houses, insurance companies, and other financial institutions. The specification of FNA services, their performance, and the terms of contract under which they will be supplied are now being drawn up by workgroups of FNA members, with the aim of providing service consistency around the world, regardless of which PTT provides it.

Two FNA members — France Telecom and Deutsche Bundespost Telekom — are involved in discussions to form yet another venture to offer outsourcing services over a European backbone. The two PTTs announced this service, called Eunetcom, in March 1992. Eunetcom plans to build a pan-European backbone, but an equipment supplier hasn't been chosen and few details are available. Still, Eunetcom hopes to start carrying live traffic on its future network by late this year or early 1994.

France Telecom has one more iron in the fire: Transpac, its packet-switching subsidiary, is inviting equipment bids for a European backbone based on ATM technology.

Meanwhile, PTTs are working together to develop centrally managed cross-border backbones for internal use, to help improve existing services and to launch new ones. These include the Global European Network (GEN), an effort involving PTTs in France, Germany, Italy, Spain, and the UK. GEN's backbone is already in place and is undergoing trials, and seven more PTTs have applied to participate in the project. A similar collaboration, called the Global Network Project, aims to have a SDH backbone in service by the end of the year. Plans for another SDH backbone have been agreed to in principle by 26 European PTTs, but this project, the Managed European Transmission Network (Metran), is unlikely to go into operation before 1996.

Telecom Denmark has adopted a different approach to collaboration. It has set up a joint venture with Maersk Data, called Temanet A/S, and it will handle almost all of Telecom Denmark's subscribers for international networking. In this low-cost approach Maersk operates the network, and provides technical support and overseas offices for Temanet. Telecom Denmark provides the domestic data network and its cables to Russia and Eastern Europe.

Peter Heywood

The Global Carrier Business Shapes Up

CARRIER OR COLLABORATION	PARTNERSHIPS	NETWORK PLANS	SERVICE PLANS	COMMENTS
AT&T	Subsidiaries, including AT&T Istel (UK), AT&T JENS (Japan), and NCR	Building European backbone through AT&T Istel	Clear channels, packet switching, frame relay, messaging	Backbone is based on multiple 2Mbps circuits
BT	Subsidiaries, including Syncordia Corporation and BT North America. New partnership with MCI Communications	Building global backbone with nodes in Atlanta, Amsterdam, Chicago, Frankfurt, Hong Kong, London, Los Angeles, New York, Paris, Sydney, Toronto and Zurich (see also MCI)	Clear channels, virtual private networks, outsourcing	Recently applied to the US FCC for a licence to offer international virtual services
Cable & Wireless (C&W)	Owns, wholly or partially, public network operators in more than 40 countries, including Australia, Hong Kong, Japan, Sweden, UK, and US	Global Digital Highway comprises fibre networks in individual countries, linked by undersea fibre cables; includes European backbone	Global Managed Data Service handles design, implementation and operation of private networks	Works with partners to offer outsourcing
Eunetcom B.V.	Collaboration between Deutsche Bundespost Telekom and France Telecom	Plans European backbone	To be determined	Announced in March 1992
Financial Network Association (FNA)	Collaboration headed by MCI and including Belgacom, France Telecom, Hongkong Telecom, Italcable, KDD, Mercury Communications, Singapore Telecom, Stentor, Telecom Australia, and Telefonica	Global backbone planned, with trials slated for June 1993	Leased lines initially; other services under discussion	Most members are Infonet shareholders
Global European Network (GEN)	Collaboration involving BT, Deutsche Bundespost Telekom, France Telecom, STET (Italy), and Telefonica	Backbone between member countries; now under trial	Not marketed directly to end-users	Other PTTs have applied to join
Global Network Project (GNP)	Collaboration involving AT&T, BT, Deutsche Bundespost Telekom, France Telecom, Kokusai Denshin Denwa, and Telstra	Backbone between members scheduled for operation by end of 1993	Not marketed directly to end-users	Other PTTs may join
Managed European Transport Network (Metran)	Collaboration involving 26 European PTTs	Plans Synchronous Digital Hierarchy backbone, for possible operation in 1996	Not marketed directly to end-users	Agreement in principle only
MCI Communications Corp	Owns 25% of VAN provider Infonet Services; founded the FNA; new partnership with BT	Link with UK; to build global intelligent network in conjunction with BT	Managed Service Links, comprising clear channels to UK, launched in February 1993	Cooperates rather than competes with PTTs
Sprint International	Interconnect agreements with public service network operators using Sprint equipment for packet-switching networks	Links 95 cities in 31 countries, via its own network and those of its partners; plans 500 points of presence outside US by 1995	Clear channels, packet switching, frame relay, messaging	Distribution agreement with Unisource in the US under negotiation
Temanet A/S	Collaboration between Telecom Denmark and Maersk Data A/S	Upgrading existing 35-node network	Mainly leased lines at present	Handles most of Denmark's international leased-line customers
Unisource N.V.	Joint venture involving PTT Telecom (Netherlands) and Swedish Telecom	European backbone in Belgium, France, Germany, the Netherlands, Spain, Sweden, Switzerland, and the UK; negotiating interconnection with other networks, notably in Japan and US	Leased lines, data services, messaging, network management; plans outsourcing	Shareholder's data service divisions turned into Unisource subsidiaries; Swiss PTT plans to join pending government approval
WorldSource	Joint venture between AT&T, Kokusai Denshin Denwa and Singapore Telecom	Global backbone planned; US/Asia Pacific in operation by end 1993, Europe by 1994	Leased lines, virtual private voice/data network, frame relay	Telstra, Korea Telecom and Unitel Canada have announced they will join

HP LanProbe II brings a new level of insight to HP OpenView.

HP LanProbe II can be managed from either an HP-UX OpenView workstation or a Microsoft® Windows™ PC.

Use HP LanProbe II with the HP OpenView Probe Manager, and manage faults more effectively.

HP LanProbe II is a second-generation Ethernet LAN monitor that spots and diagnoses network problems quickly. Using HP's newest OpenView application, HP Probe Manager, you can access the power of HP LanProbe II from an OpenView HP-UX workstation and get the most comprehensive network management from a single source.

HP LanProbe II supports RMON-MIB,* the accepted standard for SNMP network monitoring. So you can monitor all the nodes on your network - even if they come from different vendors - using the largest variety of management tools.

These tools include current statistics for fault management, and long-term trend graphing for performance tuning and network planning. Plus, you get remote packet capture and decoding for protocols such as TCP/IP, Novell, NetWare, Banyan VINES, DECnet, AppleTalk, OSI, 3COM, and XNS to help you locate problems fast and take corrective action directly.

And since HP LanProbe II complies with RMON MIB and other industry standards, your investment in probes and applications is protected.

So, if you'd like more insight into the performance of your network, call our customer information centre on 008 033 821 or Melbourne 272 2555 and we'll send you a data sheet that explains how HP LanProbe II helps you take network management to a new level.

*Remote Network Monitoring Management Information Base.

A Better Way.

HEWLETT PACKARD

Intel's Pentium: Built for Speed and Safety

The new chip's performance will grab power users' attention, but its compatibility with existing software will attract network managers.

If Intel's new Pentium processor were an automobile, it would have an Offenhauser turbine, dual overhead cams, and a V-12 engine with enough horsepower to get airborne. It also would have state-of-the-art antilock brakes, built-in child safety seats, and airbags strong enough to stop the Australian Wallabies Rugby Union scrum. Power-driven users are sure to be attracted by Pentium's performance numbers, but for network managers, the chip's safety features — specifically, its compatibility with existing PC software — may prove to be the more powerful lure.

On the performance side, Intel's latest-generation microprocessor packs about twice as much processing pop as the 80486 chip and sports a 64-bit-wide bus for faster data exchanges with peripherals and high speed networks. In terms of raw horsepower, the superscalar chip matches the best of today's RISC processors for handling distributed applications and calculation-intensive tasks, such as computer-aided design and manufacturing.

What sets Pentium apart from its RISC rivals, however, is its compatibility with the huge installed base of 80X86 software. For all its power, RISC technology hasn't found its way into mainstream network servers, primarily because of a dearth of compatible software.

With Pentium, network administrators will be able to run existing applications without a hitch, although older programs will not gain full benefit of the chip's power. For that, code has to be recompiled to take advantage of Pentium's design. Pentium also provides some features, not found in many RISC chips, that are crucial to distributed computing, such as built-in error detection for multiprocessor servers.

Intel announced Pentium in late March and already has started shipping some chips. By year's end, computer makers are expected to begin rolling out high-end servers based on one or more Pentium processors. Among the vendors now releasing or readying Pentium-based servers are systems giants like IBM, Digital Equipment Corporation, Hewlett-Packard, NCR, and Unisys, as well as PC vendors like AST Research, Compaq Computer, Netframe Systems, and Tricord Systems.

On the software side, progress probably will be slower, even though Intel has taken steps to avoid the software performance gap that continues to plague 386 and 486 systems — DOS and Windows are 16-bit systems, while the 386 and 486 are 32-bit processors. Intel has shared advance design information about Pentium with several vendors of compiler programs in the hope that Pentium compilers will be available as the first Pentium platforms are shipped. Compiler vendors working with Intel include IBM and Microsoft, along with Borland International, Metaware, Micro Focus, The Santa Cruz Operation, Novell subsidiary Unix Systems Laboratories, and Watcom International.

The Software Lag

Even with this step, some lag between hardware and software availability is inevitable. At the time of going to press, none of the compiler vendors were being forthcoming with specific dates for product shipments. And once compilers are available, it's up to software vendors to recompile their applications to run on Pentium's platforms. Software houses aren't likely to take this step until a significant number of Pentium servers are in place, a development that Intel doesn't see happening until the end of next year.

Until optimised software becomes available, Pentium users aren't likely to see much of a performance gain. Existing PC software can run on Pentium machines, but only 32-bit operating systems and applications written and compiled specifically for Pentium will be able to take full advantage of its power. Operating systems and applications optimised for Pentium will run up to 30% faster than unoptimised software, according to Ron Curry, Intel's Product Line Manager for Pentium.

Right now, however, operating system vendors including IBM, Microsoft, and Novell are saying only that they are studying ways to enhance their existing products to run best with the new chip. Microsoft's strategy is likely to focus on the forthcoming Windows NT operating system, which Microsoft has targeted to replace DOS and Windows for next-generation PCs. Windows NT is supposed to be platform-independent, which means it will run on high performance RISC processors as well as on Intel PC chips. For the first time, therefore, network managers will have the option of

Figure 1: The New Chip Inside

The key to Pentium's performance advantage is its use of two separate processing elements (integer arithmetic and logic units), which enable the processor to execute two instructions simultaneously.

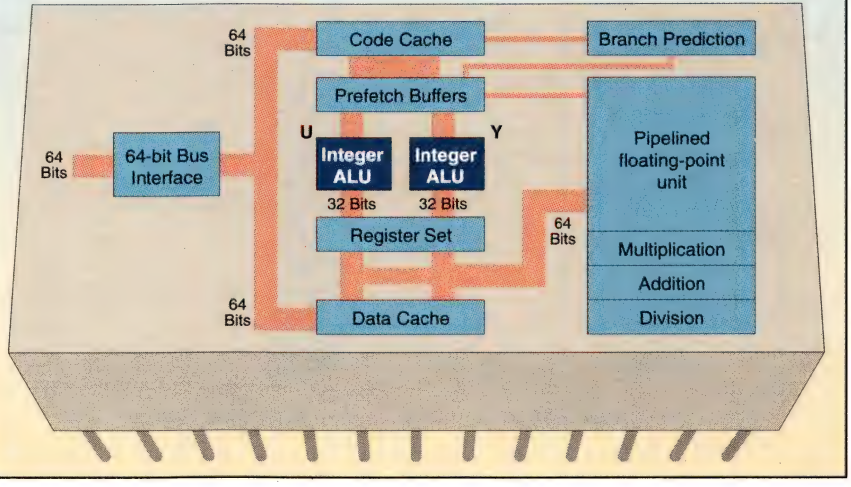

Market Watch

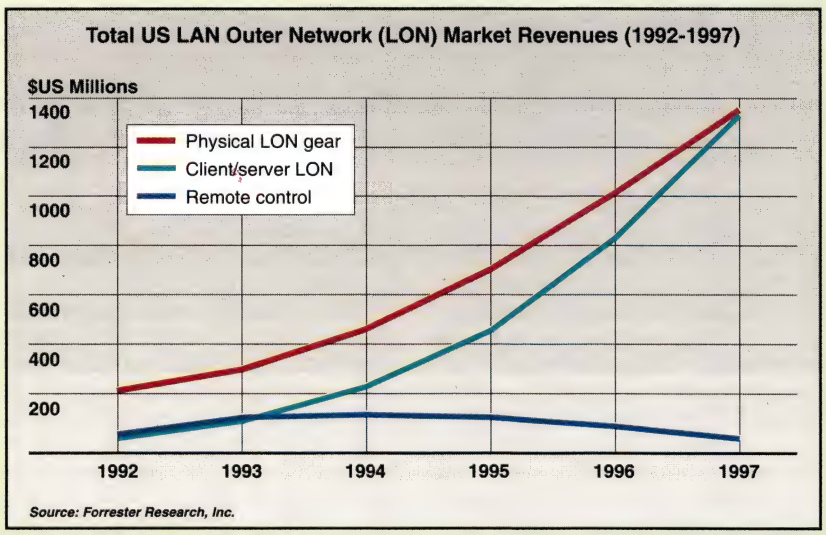

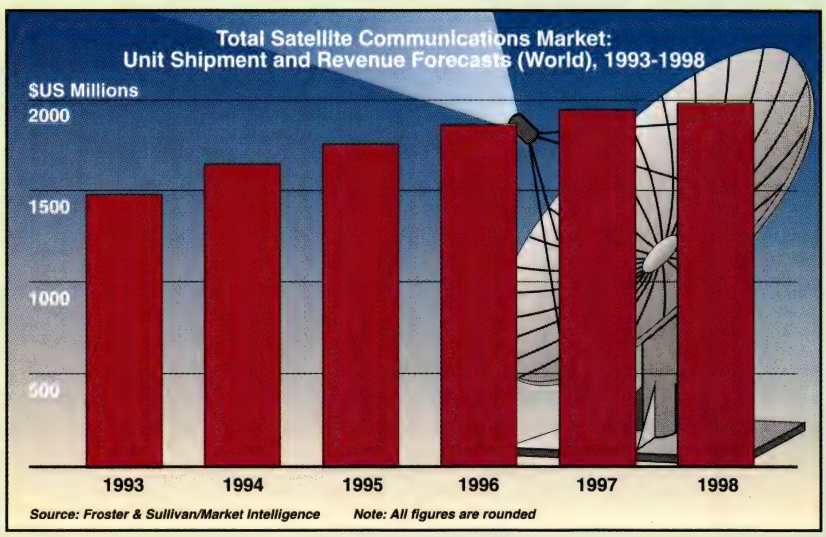

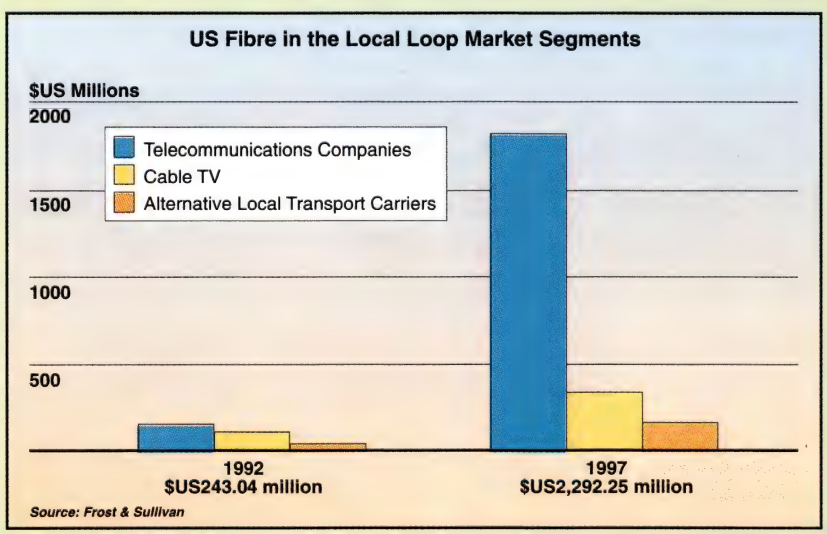

choosing non-Intel hardware to run distributed applications on corporate LANs.

For Netware users, performance improvements from Pentium will come from distributed applications, such as database engines and transaction processing. Pentium is unlikely to be needed for NetWare servers employed exclusively for file, print, or communication services. Input/output speed, rather than raw processing power, is the primary bottleneck in such configurations.

The Pentium Difference

Although Pentium is an extension of Intel's 80X86 line, the new chip represents an important change in design approach. The most critical is Intel's use of a superscalar architecture, a common feature for RISC processors but one that Intel has never used in its CISC designs. Pentium's superscalar pipeline features two processing elements — called integer arithmetic and logic units (ALUs) — that can execute separate instructions simultaneously (see Figure 1 on page 29). In addition to separate processing channels (Intel calls them U and V channels), Pentium also has separate code and data caches, a redesigned floating-point unit, and a 64-bit external bus, which will allow faster data exchanges than the 32-bit buses of existing 80X86 processors.

Each of Pentium's integer ALUs offers power on a par with an individual 80486. The ability to incorporate two processors into one chip comes as a result of improvements in chip manufacturing techniques. Pentium contains about 3.1 million transistors, more than twice as many as the similarly sized 486 holds.

Since the processing of distributed applications involves integer calculations almost exclusively, the use of two separate integer ALUs should double the speed with which applications are processed — in theory, at least. Software has to be written and compiled to make full use of the two processors, a development that, as noted, will take some time.

For high-end applications that involve much heavy-duty number crunching, such as CAD/CAM and forthcoming multimedia applications, a chip's ability to handle the floating-point calculations is crucial. The floating-point unit (FPU) built into Pentium is radically different from that used in previous Intel chips. Pentium's FPU contains circuitry which is fully dedicated to specific arithmetic operations: multiplication, division, and addition. Again, the ability to incorporate more transistors on a single chip has enabled Intel to improve floating-point operations.

Pentium's 64-bit data bus will allow data to be moved from external caches and peripherals faster. Only one other major vendor now offers a processor with a 64-bit bus: DEC, which sells the Alpha line of RISC chips.

The 64-bit bus has special significance for network managers. It means adaptor vendors will be able to produce 64-bit cards. These will be invaluable for high speed LANs, such as FDDI or ATM networks. Mass storage systems found in servers, such as RAID (redundant array of inexpensive disks) systems and SCSI (small computer systems interface) drives, also will benefit from the wider bus.

A Faster Local Bus

With Pentium Intel also announced the 82430 PCIset chip set, which forms the basis for a high-performance local bus. The PCI local bus supports 64-bit data transfers and multiple bus masters, which means that add-in cards plugged into the local bus can take control of the bus and transfer data without the intervention of the Pentium processor.

Graphics boards are the most likely candidates for use with the local buses. However, SCSI storage systems and network adaptors also will run faster with the PCI local bus. Vendors of these products have yet to make specific announcements.

One potential snag with the PCI local bus is that it is not compatible with the local bus found in 486-based servers. These implementations are based on the VL-Bus standard created by the US Video Electronics Standards Association. Add-in boards that work with a VL-Bus thus will not work over a PCI local bus.

Intel's Curry says the chip maker didn't adopt the VL-Bus because it cannot handle multiple bus masters and 64-bit data transfers. But VESA counters that it will enhance its bus design to match the abilities of the PCI chip set. If this happens, systems makers could build Pentium-based computers with VL-Buses in addition to PCI buses, as well as conventional EISA or MicroChannel buses, for that matter. This would allow existing add-in cards to be used with Pentium machines.

Because Pentium is targeted specifically to server applications, Intel outfitted the chip with several features that are becoming prerequisites in LAN servers. One is a mechanism to ensure that data in all server caches is consistent. All multiprocessor systems must have cache consistency: If data is changed in one cache, appropriate changes must be made to related data in another cache.

Pentium implements the MESI (for multiplied, exclusive, shared, invalid) protocol to ensure cache consistency. In the past, the functions performed by MESI had to be implemented in software, exacting a performance penalty.

With input specifically from Novell, among other vendors, Intel also built new data integrity features into Pentium. Parity checks are conducted on all major onboard elements, including caches and registers, to detect any errors. The chip also performs functional redundancy checking, making it easier for two Pentium processors in a single system to run at the same speed and work on the same task.

Figure 2: RISC-Free Number Crunching

In Spec92 benchmark tests of integer calculation performance conducted by Intel, the Pentium processor doubled performance on Intel's 486 chip and outperformed all RISC competitors except DEC's Alpha 21064, a 133 MHz processor.

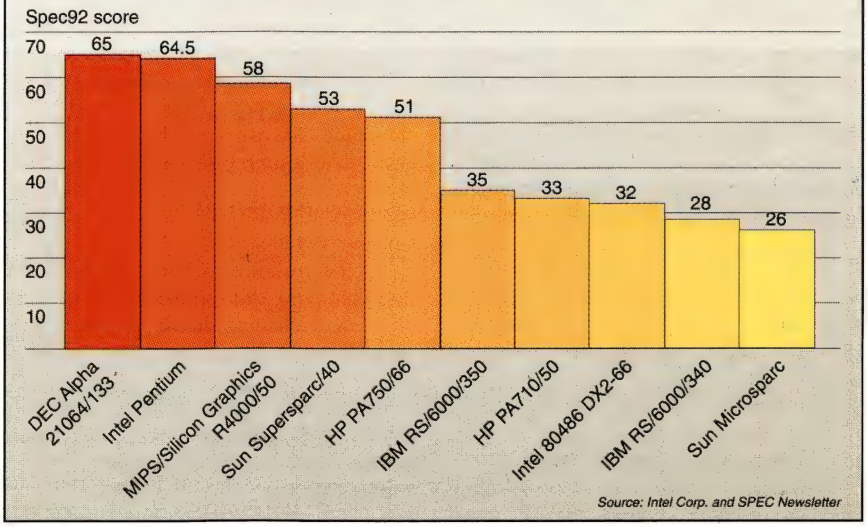

Features like cache consistency and error checking are icing on the performance cake. Initial Pentium chips, which will be available in 60 and 66MHz varieties, will now provide processing power significantly greater than 80486s and matching or exceeding that of current RISC processors, while maintaining code compatibility with the 80X86 line.

Pentium's performance possibilities are hinted at in the results of a Spec92 benchmark of integer performance, conducted by Intel and published last December in the *SPEC Newsletter*. Intel's results show Pentium performing almost twice as fast as an 80486, and faster than all other RISC chips except DEC's 133MHz Alpha 21064 (see Figure 2).

The Prognosis

Spec92 is just a measure of raw horsepower — how Pentium performs in finished systems remains to be seen. One thing all vendors agree on is that Pentium will not deliver optimum performance if hardware vendors simply substitute it for 80486s. New hardware designs, including faster memory, faster buses, and new peripheral components such as interrupt controllers are essential.

New designs also will be crucial to ensure heat dissipation. Pentium, with its higher clock speed and larger number of components, will run hotter than existing 80X86 chips. Systems designed for the 80486 most likely will not have an adequate cooling system for Pentium. Peripherals built to take advantage of the chip also will require extra cooling.

The first Pentium implementations likely to be available will be 486 machines designed to be upgraded to Pentium. The new chip has more pins than the 486, but Intel had released preliminary design information so that computer makers could include the necessary connectors. Vendors with Pentium-ready machines also have built extra cooling capacity into their boxes.

Another way to build Pentium into existing systems is to substitute add-in cards or daughterboards containing Pentium for similar boards containing the 80486 in systems designed with Pentium in mind. With this approach, systems makers could provide the enhanced caches and other components that allow Pentium to be used more efficiently. Most vendors that offer board swapping have built in the necessary cooling capacity.

IBM's high-end Model 95 multiprocessor servers are designed with 80486s on removable cards, as are NCR's System 3000 computers. Both companies say they will provide Pentium upgrades for their servers. Netframe Systems, the first vendor to offer PC LAN superservers, says it will upgrade to Pentium in this manner. Its C/S (Client/Server) server line holds from one to four 80486 processors. Each processor comes on its own add-in card. The computers will accept similar numbers of Pentium add-in processors in the future.

Systems built specifically for Pentium processors probably will hit the streets next year. Eventually, Intel will ramp up Pentium's clock speed to 100MHz. Low-power versions for laptops also should start showing up next year.

Eric Hindin

High Speed Networking

Can Fibre Channel Challenge ATM?

In this age of ATM and FDDI, channel-attached architecture might seem like an odd choice for connecting high speed LANs. But a channel design that's usually associated with host-peripheral connections could emerge as a faster, cheaper alternative to today's high-end LAN architectures.

Strange as that may sound, the standard — called Fibre Channel — already has garnered substantial support among industry heavyweights such as Hewlett-Packard and IBM. They reckon that the ultra-high speed channel architecture (it operates at up to 1Gbps) can be converted for LAN use by adding a switch, specified in the new standard, that handles multipoint addressing.

However, to gain mainstream acceptance as a LAN technology, Fibre Channel will first have to beat out ATM (asynchronous transfer mode). Given the networking industry's overwhelming support for ATM, that's about as likely as Australia winning next year's soccer World Cup. Almost two dozen vendors are readying ATM switches for market, and 1994 is likely to be a better year for ATM rollouts.

ATM products are sure to limit demand for Fibre Channel. "We're not talking about the next Ethernet here," according to Richard Villars, Director of network architectures at market researcher International Data Group. Nevertheless, Fibre Channel has a card up its sleeve that could guarantee it respectable market share: it provides much lower network delay (latency) than ATM for some data-intensive LAN applications (see the table below).

Stable Standard

Fibre Channel is being developed by the X3T9.3 committee of the American National Standards Institute (ANSI). Most of the standard is now stable, and it is scheduled for final ratification in the third quarter of this year.

Fibre Channel targets three types of connections: conventional point-to-point I/O applications, like mass data storage; clusters of high-end workstations; and high speed switched LAN connections.

According to the Fibre Channel Association, a vendor group, three United States companies are working on Fibre Channel network adaptors: Emulex, IBM, and Interphase. A fourth vendor, Canada-based Canstar Communications, is working on a Fibre Channel switch.

In the meantime, Ancor Communications plans to ship both a Fibre Channel switch and adaptors shortly. Sun Microsystems also is working on an Sbus adaptor for use in storage applications. Additionally, Hewlett-Packard is known to be working on Fibre Channel products, but declines to give details. At least 15 other companies are now working on silicon or fibre optic components for Fibre Channel products, according to the FCA.

The extremely high speeds supported by Fibre Channel — up to 100 times faster than existing Ethernet LANs — make it an obvious candidate for applications like mass data storage.

But vendors of Fibre Channel products think such performance also helps open up a whole new range of possibilities for distributed computing among local groups of high-end workstations. "The high speed networking capability provided by Fibre Channel will enable clustered workstations to address whole new classes of data-intensive applications with a level of performance once achievable only on a supercomputer," says Phil Hester, IBM's advanced workstation division Vice-President for systems technology.

Both IBM and Hewlett-Packard are working on cooperative operating system software that will enable multiple workstations to process tasks as a single device. IBM is endorsing Fibre Channel as the most suitable technology for clustering workstations, including its RISC System/6000.

While Fibre Channel offers clear benefits in linking clusters of high-end workstations, today's PCs aren't fast enough to drive demand for Fibre Channel as a general-purpose LAN technology — the third application for which the standard is targeted. Nevertheless, predicted exponential increases in personal computer performance mean PCs and applications requiring Fibre Channel service "could appear within three years," says Joe Mathis, principal senior technical staff member at IBM and technical editor of the ANSI architectural document for Fibre Channel.

Notably, not all of the vendors now pushing the Fibre Channel standard are touting it for use on LANs. Sun's perspective "is that we are interested in Fibre Channel as an I/O technology, not as a local area network technology," says Anil Uberio, Group Marketing Manager for networking products at Sun.

Changing Channels

Fibre Channel employs the same channel-based architecture as network adaptors used for host-peripheral connections or point-to-point data storage. In general, channel-based adaptors carry data packets at the highest speed, and with the least network delay (latency) possible.

To achieve high speeds and low delays, channel-based adaptors do most processing in hardware. LAN adaptors, in contrast, usually perform most processing in software and consequently run at slower rates. "The software [in channel adaptors] is limited to one very simple driver," says Wayne Rickard, Principal Engineer at Emulex. In LAN adaptors, a much larger driver is used to handle a variety of functions, including handshaking, address selection, and error correction, he says.

Conversely, the pared-down design of channel products — which perform only elementary error correction and support no sophisticated station management or addressing facilities — also makes them inherently ill-suited for LAN service. The Fibre Channel standard gets around this limitation by defining a switch that establishes multiple simultaneous point-to-point connections among adaptors connected in a star topology.

The Fibre Channel spec also defines an encapsulation technique to handle both high-performance channel protocols and LAN protocols. As LAN and channel traffic passes onto the Fibre Channel network, it is placed into Fibre Channel 'frames,' as they are referred to in the standard. On the other side of the Fibre Channel network, the frame is stripped off. Interfaces supported by Fibre Channel include HIPPI (high-performance parallel interface), IPI (intelligent peripheral interface), and SCSI (small com-

Fibre Channel vs ATM

	FIBRE CHANNEL	ATM
Original Use	Channel-attached connections	Telecommunications
Speed	133Mbps, 266Mbps, 530Mbps, 1Gbps	155Mbps, 622Mbps
Packet Format	2K frame, 24- to 28-byte header	53-byte cell, 5-byte header
WAN Capabilities	No	Yes
Congestion Control	Hardware-based; frames sent only when space is available	Software-based; cells discarded when network becomes congested
Support for Time-Sensitive Traffic	Via dedicated connections; switched capability planned	Via switched connections

puter systems interface). Fibre Channel also will support ATM, Ethernet, FDDI, and Token Ring LAN traffic.

The elementary design of channel adaptors has led ANSI to pick a highly switch-centric addressing scheme for Fibre Channel. The adaptors simply add an address header to each frame. The switches provide all the intelligence necessary to route the frames across the network, including error recovery.

To make the standard relevant to as many installations and applications as possible, Fibre Channel allows product developers to choose among a wide variety of speeds (133Mbps, 266Mbps, 530Mbps, and 1Gbps) and cabling types (multimode and monomode fibre, thick and thin coaxial, and shielded twisted-pair [STP] wiring).

Transmission distances vary depending on the combination of speed and media. For example, the standard defines a maximum distance of 25 metres for 1Gbps traffic over thick coaxial cable and 100 metres for 133-Mbps traffic. However, the standard recommends that STP cabling not be used at speeds above 266Mbps.

To transmit frames, Fibre Channel specifies an IBM-developed encoding scheme called 8B/10B. For network management, the X3T9.3 committee has under development a new management information base (MIB) that will enable Fibre Channel products to be monitored from any SNMP- compliant console.

Priority Service

To ensure efficient transmission of both channel and LAN traffic, Fibre Channel defines three classes of service. Users select service classes based on the characteristics of their applications, like packet length and transmission duration, and allocate the services by configuring the Fibre Channel switch.

Class 1 service configures Fibre Channel switches so that a dedicated channel is set up between two network nodes — in effect providing the equivalent of a dedicated physical connection. During call set-up, frames are sent to notify switches along the route that the link should be made unavailable to other devices. Once the link has been established, switches do not have to look up the frame's address in their routing tables. Because Class 1 provides dedicated bandwidth, "it is the best bet for sustained, high-throughput transactions such as real-time graphics and mass storage," claims Kumar Malavalli, Director of Architecture and Product Planning at Canstar.

Class 2 service allows bandwidth to be shared by multiplexing frames from multiple nodes onto the same channel or channels. In this mode, address headers are compared with those in a routing table at each

Continued on page 44

Technology Update

Work Begins on ATM and Sonet MIB

The Internet Engineering Task Force (IETF) has chartered a working group to construct an SNMP management information base to define managed objects for asynchronous transfer mode and Synchronous Optical Network systems. According to the IETF the Atom MIB will include ways of describing the equipment, interfaces, networks and services in ATM and Sonet networks to standard SNMP management stations. The group has said it will try to coordinate its work with that of other IETF groups and the ATM Forum. The Sonet definitions will be aligned with the SD1/DS3 MIB, which already has been defined by the IETF. The Atom MIB group hopes to have a proposed standard by the end of the year.

Government Approves RSA Public Key Encryption

The Australian Department of Defence has just approved in principle the use of RSA public key encryption technology for use by Federal Government departments. The technology has been approved for applications including digital signatures, document authentication and key management, but has not been approved for general data encryption. A spokesperson for the Department said the decision to approve the technology had been taken because conventional systems were becoming increasingly unable to cope with the widespread demand for document security systems due to growth in the use of computer networks. The RSA encryption system was invented by Professors Rivest, Shamir and Adleman at MIT in 1978, and has gained wide acceptance, being approved as a standard by CCITT, EDIFACT and the ISO. A number of major vendors have licensed the technology for inclusion in products, including Apple, AT&T, IBM, Microsoft, Novell, Lotus and Sun.

MNP 10 Granted a Patent

US-based Microcom has been issued a patent for Microcom Networking Protocol 10 (MNP 10), which improves data transmission performance by adapting the size of data packets to line conditions. Using MNP 10, two modems exchange information regarding line conditions using 'Link Management Idle' packets, which are sent during pauses in the transmission of user data. Currently, modems use relatively large data packets (128-256 bytes) at the beginning of a session, reducing the packet size if too many errors occur. With MNP 10, the modem begins by sending smaller packets, and then increases the packet size as conditions permit. Microcom says the advantage of this approach is that it allows modems to respond to dynamically changing line conditions, such as a connection that improves during the course of a call, increasing overall throughput. Several manufacturers now license MNP 10 technology, including Compaq, AT&T Microelectronics and Rockwell International.

Streamlining X.400 Addressing

A newly-formed implementors group is aiming to build a cohesive X.400 network in the US. The US National Message Transfer Services (US-NMTS) group is an ad hoc organisation consisting of eight major X.400 service providers and several companies that are X.400 users. The aim of the group is to create an X.400 messaging infrastructure based on a single administrative management domain (ADMD), which will allow each user to maintain one X.400 address across multiple public X.400 carriers. Under the current system, each carrier issues its own X.400 address, which contains fields that specify the user's address and the route by which a message must travel. This can be cumbersome for users who need connectivity to more than one X.400 service. The first issue the group will attempt to resolve is the setting up of an addressing scheme similar to that used for long-distance phone calls.

Modem Management MIB

A working group will be formed to define an SNMP MIB for modems. Modem makers have been using proprietary protocols for management, but an informal session held at an IETF meeting in March attracted a large number of modem vendors keen to establish a common set of SNMP management definitions. The definitions could monitor functions like a modem's line and DTE interfaces, its hardware and software parameters, error conditions and the number of calls made. SNMP could also be used to configure and reconfigure modems.

APPC Users Get 3270 Access

APPC developers have proposed a new utility, called APPC3270, that will allow 3270 traffic to run over APPC sessions, freeing up workstation memory. The utility will mean that network managers using IBM's Advanced Program-to-Program Communications (APPC) to develop client-server applications will no longer have to load two protocol stacks — LU 6.2 for APPC applications and LU 2 for traffic from existing 3270 terminals. APPC3270 will be comparable to TCP/IP's TN3270 utility in that it uses an emulation program to access host applications. The APPC3270 utility will run either on a gateway or at the host, mapping or converting the APPC data to LU 2 format. Using the APPC3270 utility, users will be able to access both APPC client-server applications and host applications by loading a single APPC protocol stack.

ATM

QPSX Unveils ATM Chip Set

Australia's world-leading telecommunications innovator, QPSX Communications, has taken another leap into the future with the development of an ATM (asynchronous transfer mode) chip set capable of bringing fast-packet switching technology to the desktop.

This is a silicon translation (and improvement) of the techniques used for Fastpac-type public metropolitan area networks (MANs), which was also chosen by the IEEE as the basis for the international DQDB 802.6 MAN standard.

DQDB (distributed queue dual bus) is the switching and access technology widely used in the US and Europe for SMDS (switched multimegabit data services) for linking LAN-to-LAN across a city, or more recently across a nation, and it is now about to be used by the British as feed service for their new data 'super-highways' research network, SuperJanet.

The key component in their new chip set is an ATM cell-processor which can handle data at rates up to 622Mbps, and handle different interface protocols; each chip has about half-a-million circuits, QPSX says. The data-rate being handled is the first commercial emergence of the STM-4 (Synchronous Transfer Module) standard which eventually will be used on Broadband ISDN. The chip set also handles the slower STM-1 (155Mbps), and STS-1 (51Mbps) physical layer standards. The physical layer protocols for STM standard frames include monitoring and management signals so that a separate management network is not needed — and this makes the chip set ideal for private networking as distinct from the more common use of DQDB technology in the public network.

There are now three chips in the set. One is specific to QPSX products, and the other two are for third-party products. Of these, the queue/cell processor is microcoded to allow the same silicon to be used to perform a variety of different functions. So the same chips can also be applied to pure ATM point-to-point links or to shared-media DQDB-type MANs, either with the famous queuing access method in operation or with pre-arbitrated links having guaranteed access. The queuing protocol allows for multiple levels of access priority.

The QPSX chip set generates the required physical layer framing (at the various rates and standards), and performs stuffing and timing insertions to maintain network synchronisation. It also handles the segmentation and reassembly of large packets to create the ATM/DQDB cells, and provide the ATM Application Layer functions for Classes 3, 4 and 5 (AAL 3/4 and AAL5). By providing standard AAL packet formats, the set can map to both ATM cells and 802.6 MAN slots.

Service-specific functions such as address validation, access classing, and charging are also included. The chips can map VPI and VCI — the standard 'virtual path indicator' and 'virtual circuit indicator' addressing standards of ATM.

More MAN Success

Back on the export-front (with the earlier MAN developments) QPSX has also had a number of recent successes in Europe and America. MCI's new HyperStream national public fast-packet data service has now gone online across America, and is expected to grow quickly to provide 300 access points by mid-1993. In Germany, the Deutsche Bundesposte Telekom Datex-M service will reach into ten major cities by next year.

And in the United Kingdom, BT's first 'data super-highway' called SuperJanet (for Joint Academic Network) will eventually link 50,000 terminals (including DQDB local networks) in 200 universities and research institutions around the country. This is an £18 million project spread over four years, using QPSX and SDH (synchronous digital hierarchy) technologies, with data-rates initially from 34Mbps to 140Mbps, rising later to the full 622Mbps when requirements and funding coincide.

Stewart Fist

Modems

Modem in the Middle: Sizing Up V.32terbo

Picking a next-generation modem just got more difficult. A group of modem and chip makers, dissatisfied with the CCITT's slow progress in defining the V.fast specification, are now developing V.32terbo, an alternative spec that they say will bring products to market more rapidly.

V.32terbo defines data rates that are up to a third faster than V.32*bis*, the current high speed specification. In addition, V.32terbo modems are likely to be cheaper than those based on V.fast because the technology is less complex.

But leading V.fast backers say V.32terbo has enough technical shortcomings to justify waiting for the new CCITT specification. The clearest sign of a rift occurred in February, when six key modem makers called a press conference at the Communications Networks (Comnet) trade show in Washington, D.C., to announce their opposition to V.32terbo.

For network managers, the controversy means greater uncertainty over shipping dates, feature sets, and even data rates. More significantly, the divergence could break apart the once ironclad link between speed and standardisation.

Implementors' Agreement

Despite V.32terbo's name, which is deliberately reminiscent of the CCITT's V-series specifications, the spec is actually an implementors' agreement rather than an official standard.

In January, a group of 18 modem and chip manufacturers banded together and agreed to support V.32terbo, which modifies the existing V.32*bis* recommendation to run at 16.8Kbps and 19.2Kbps instead of its current top speed of 14.4Kbps (see table on page 36). In addition, the group plans to pursue efforts at formal standardisation of the spec, although no standards body has been chosen. V.32terbo backers include AT&T Paradyne, Data Race, Multi-Tech Systems, National Semiconductor, and Penril Datacomm Networks.

V.fast, which defines a top speed of 28.8-Kbps, is currently under development by a special committee of the CCITT (the International Telegraph and Telephone Consultative Committee). The committee expects to complete a draft specification sometime this year, with products presumably following soon after.

V.32terbo opponents believe the whole idea is unnecessary, since V.fast is forthcoming; and technically unsound, because it lacks key features for high speed data transmission. In particular, V.32terbo's detractors say modems implementing the spec won't run at the advertised data rates much of the time. The anti-V.32terbo faction at Comnet was made up of six modem maker heavyweights: General Datacomm, Hayes Microcomputer Products, Microcom, Motorola Codex, Racal-Datacom, and Rockwell International.

"We are going to have a high speed standard soon. We just don't need an in-between standard," says Vede Krishnan, Senior Fellow for access product development at Racal-Datacom. "People can wait a few months for V.fast."

But V.32terbo supporters say that's the wrong attitude. In their view, V32terbo simply offers users another data rate option — one that will be available sooner. V.32terbo "is a competitive advantage to users today," says Haig Sarkissian, Manager for data communications market development at AT&T Microelectronics. Even if the V.fast committee finishes its work in the

Jtec

Solutions

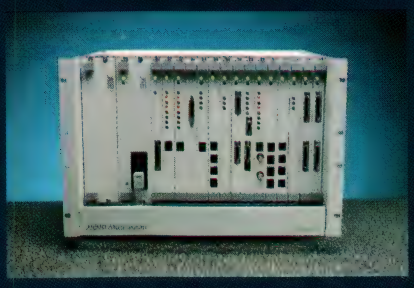

Networking Multiplexers
for
Voice
Data
Video *and*
ISDN Applications

A complete family including support for smaller sites

Data only applications supported by Advanced ISDN Terminal Adaptors

Intelligent PC Add In cards for Advanced DeskTop Applications

Comprehensive yet affordable Network Management

Head Office
Jtec Pty Limited
ACN 003 169 088
Unit 3
118 - 122 Bowden St
Meadowbank 2114
NSW Australia

Phone

International	612-809-6933
Sydney	02-809-6933
Melbourne	03-263-0200
Brisbane	07-229-9137
Canberra	06-256-3480
Europe	35-321-277-422

Fax

International	612-809-6619
Sydney	02-809-6619
Melbourne	03-263-0215
Brisbane	07-221-6625
Canberra	06-201-3003
Europe	35-321-270-420

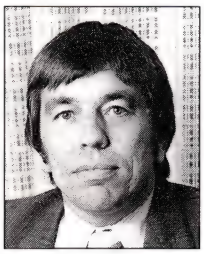

Network Design and the Integration of OSI

Alan Lloyd

As discussed in last month's *OSI Tutorial*, today's networks are moving from a terminal-based host centric environment to a message/file-based decentralised environment. Because of the need to deal with the associated network scale and migration policies, they are usually integrating proprietary, de facto and standards-based networking technologies. For network designers, the challenge is to manage this change in an integrated and cost effective manner.

Networks are in fact distributed systems which must be integrated with the strategic goals of the business. Symptomatic of today's technology push is the desire to keep on installing 'whiz bang' tools on the desktop PC. The question now asked in cost conscious business environments is what productivity gains these tools actually bring to the business. Few can answer the question.

Although network design issues are complex, the goal of my next few OSI tutorials is to provide a basic framework for addressing network design in the business environment. Networks can be designed without this consideration as 'a set of protocols' or a number of low cost communications devices and software modules (e.g. LANs and servers). While these networks usually work up to a given scale, beyond this point, technical or operational problems creep in.

The main points to be considered in the large scale network design issues revolve around carriers and purchasers. Carriers devote considerable resources to network traffic management groups and the use of international standards. These groups determine the dimensioning of the overall network so that it provides workable service levels that can be supported with formal contracts. This fundamental aspect of network traffic management seems to be ignored in some of the approaches to large scale network design being used today.

At the other end of the scale, purchasers of low cost communications devices are often buying the parts of a network from suppliers who operate on low profit margins. Without this profit margin the associated services and support will be limited or available only for an extra cost. It is then up to the purchaser to produce/commission the total system. This is like buying all the parts of a car from spare parts dealers then assembling it and making the car work. While this might be fine for a kit car, you wouldn't do this if you were building a mission critical vehicle like an ambulance.

This is one of a series of open systems tutorials by Alan Lloyd, Strategic Developments Manager for Datacraft Australia. Alan represents Australia on numerous international standards bodies and is the co-author with Gary Dickson of *Open Systems Interconnection* (Prentice Hall, 1992).

Today, the cost of IT components purchased compared with the cost of operation and integration is probably in the order of 1:10. Buying the wrong or the cheapest parts could extend this ratio dramatically. Buying low cost parts for open distributed systems is not the best way to deliver reliable, large scale IT system infrastructures.

In today's network environment, the IT manager is facing a number of major network design issues. These issues are technical, commercial and operational. From a technology perspective, the user of networking technology today needs to work out strategies for:

- Embracing different carrier communications technologies such as LANs, ISDN, X.25, DDN, Fastpac, frame relay, B-ISDN and ATM;
- Supporting the current investment in host-centric proprietary networks;
- Supporting current host-centric proprietary applications and systems;
- Integrating de facto standards like TCP/IP for Unix and distributed processing;
- Adopting OSI standards when appropriate for network scale and global internetworking;
- Dealing with the move from terminal-based networks to file/message-based networks;
- Integrating the management systems to reduce the cost of operation; and
- Adopting new communications technologies on a cost benefit basis.

From an operational perspective, the IT and network strategist needs to address:

- Downsizing — moving from or integrating central hosts and PCs, file servers and workstations;
- Standardisation — integrating technologies from a range of suppliers and utilising standardised technologies where it's possible to do so;

Figure 1: Business Approach to Distributed Systems Design

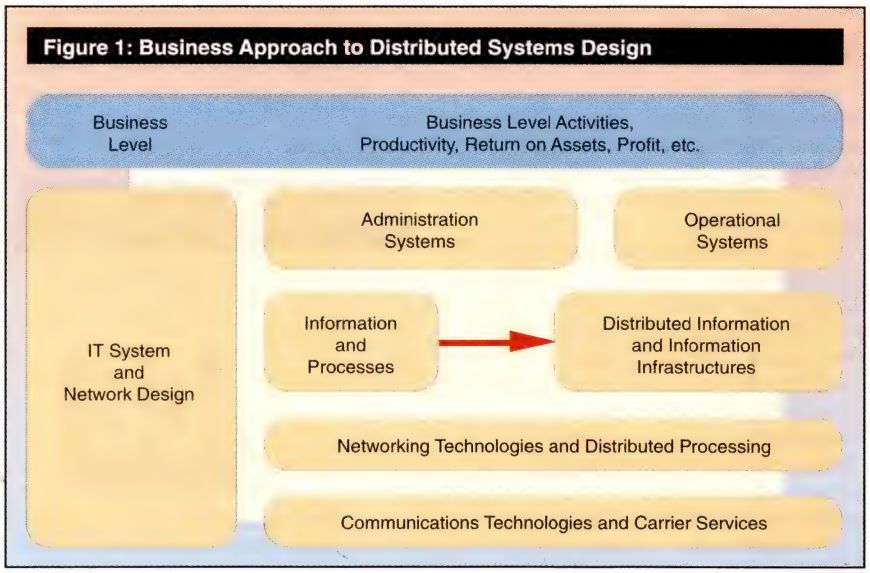

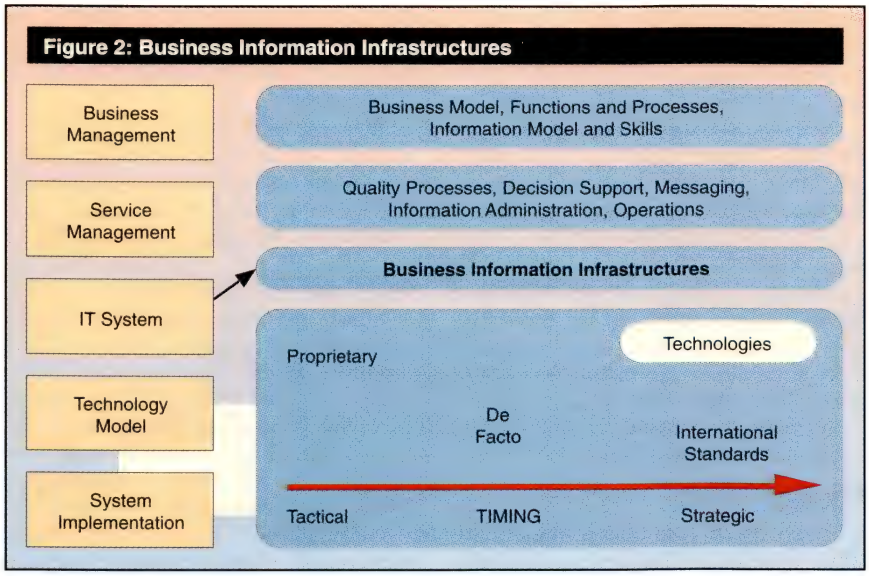

- Carrier services and costs — optimising the carrier's communications services;
- Distributed computing — positioning applications around the network, the evolution to client/server architectures and network scale;
- Operational costs — optimising the network and its operational support costs;
- User service levels — providing a respectable level of service to the user, sometimes formalised in contracts;
- Information security — determining and controlling what information is being used across the distributed system;
- Complexity — dealing with the increasing technical and business interdependencies of today's world;
- Global EDI and communications — by guaranteeing the ability to interconnect with new business opportunities as and when they occur; and
- Rate of change and return on investments — dealing with the rapid retirement of equipment because of changes in price/performance and optimisation of operational staff skills.

Design Framework

To assist in the process of system design, abstract frameworks are useful. Figure 1 identifies the network design framework from a business perspective. The lower components of the diagram identify the business information infrastructure, networking and communications, which are derived from the logical network design diagram described in last month's tutorial. The operational and administration components identify the system and network management needs from the perspective of technology, skills base and business procedures. The top level is the core business requirements.

Such a business related framework is necessary because business today is now fully supported by, and highly dependent on, its IT infrastructure. Distributed information is being used to underpin the processes, productivity and intellectual assets of the business. Because information is stored, processed, and communicated over the distributed system and that system is organised according to the business and operational aspects of the company, the network design must embrace these fundamental issues if the IT system is there to serve the business.

Procuring technology is the easiest part. The cost of operation is the key. An off-the-shelf protocol or product approach to network and system design should not be applied for the large scale business enterprise, especially when the cost of technology is falling and the cost of people and the support services rising.

Strategic Approach

The mention of applying a strategic approach to IT within a business sometimes strikes at the hallowed area of what the business is doing over the longer term and more importantly how IT might serve it. Because this is often seen as difficult, the business network is usually left to tactical growth techniques.

Taking a broad view of the business and how it applies IT is a better way of working, particularly as we move into the areas of international electronic trading. Just installing in an arbitrary way things such as FDDI, or running a range of services, giving the users what they want, spending money on the equipment and the training is really not the way to do things! With the rate of change in technology and business parameters, this approach is a high cost, low productivity, large scale minefield.

The first step with IT system development is to take at least a three year, if not a five year, view of the business. Typically, the business will have an operational model based on its location, distribution strategies and functions, such as transport, production, storage, sales and marketing. Associated with these operational functions are administrative functions — finance, human resources, quality, legal services, strategic planning, executive decision support, etc.

Growth parameters should be applied to the above according to the five year plan. From this a business model and a distributed business model can be formulated.

The next step is to build the information model, i.e. what information is needed, how it is processed (application functions), where it is needed, who owns it, who can access it and modify it, what is the lifetime of it, etc. With these parameters, the business information model will be refined to form the distributed information model. As the design starts to unfold, information processing and storage mechanisms, machines, and people will be positioned around the network. Once things are positioned, they will need names to identify them and addresses to indicate where they are.

Figure 2 shows a top down approach to the business, business information infra-

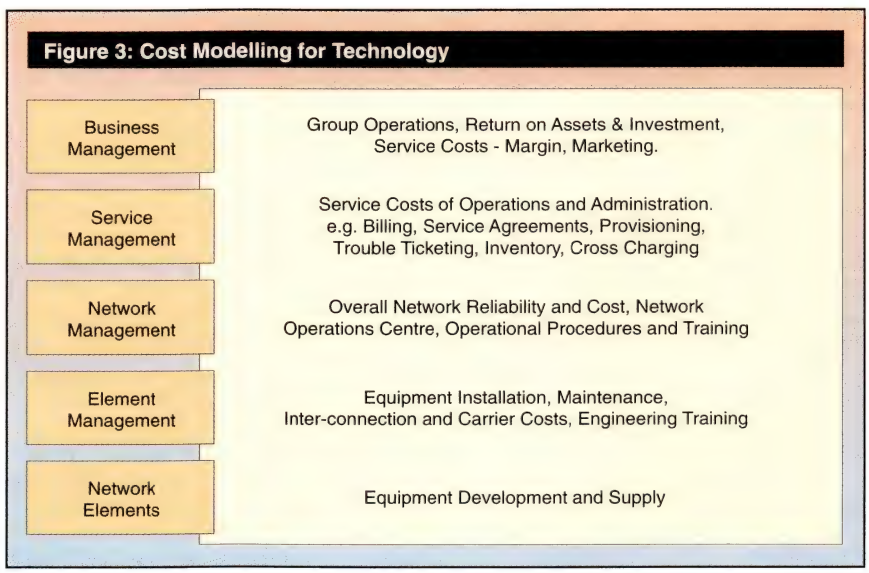

structures and the migration or integration requirements of the proprietary, de facto and OSI technologies for communications and networking. It represents the relationship of the business and its supporting technology. It also provides a general model representing how systems management and system costing can be introduced. Note that it places the business information and the procedures associated with it and its distribution as the bridge between the business itself and the network.

The resultant IT system is based on a technology model implemented over a period of time using proprietary, de facto and OSI technologies. This implies that the IT strategy, if appropriate, has to address the shift from host-centric networks to integrated, logical mesh, router-based networks and then move to use OSI technologies. This process is not a clean switch, but an evolution which integrates these technologies as appropriate. Nevertheless, the network migration issues of when, how big and how much are business related. They must be managed and resourced on that basis.

In some cases the IT departments are having to cost-justify the services they provide to the organisation based on service level agreements. This reflects the approach to running business functions as cost or profit centres i.e. business units. The IT department is then faced with providing a level of service for a given cost. However, the IT department is then exposed on two fronts. Firstly, if it does not get direction from the business level with respect to business policy, expansion, distribution, trading partners, business procedures, etc., it may not invest in an IT strategy which supports the business. Secondly, it is exposed to user demands for more features, functions and services, otherwise known as End System Anarchy. It is derived from the most commonly used and dangerous statement of 'Give the users what they want'!

Figure 3 on page 43 provides a similar abstract framework for cost modelling the IT services against the business requirements. It can be seen that the cost of the equipment is the lowest cost of the whole system, especially in a technology competitive marketplace.

In conclusion, the IT strategy which incorporates business modelling and consistent cost modelling should not be seen as a way of planning IT for the sake of procuring technology. It should be seen as a vehicle for consolidating IT systems and business directions. It will also give some measurement to the level of productivity gained from IT versus its costs. Information security and the intellectual property aspects will also be identified.

Next month's tutorial will look at messaging infrastructures and management systems from an operational perspective.

Alan Lloyd

Fibre Channel from page 33

switch on the network. Class 2 is the better choice for bursty or interactive traffic.

Both Class 1 and Class 2 send acknowledgment frames confirming frame delivery. Class 3 service is identical to Class 2, except that frame delivery is not confirmed.

The Fibre Channel standard also defines an optional mode called intermix. In this configuration, Class 1 frames are guaranteed a specified amount of bandwidth, but Class 2 and Class 3 frames are multiplexed onto the channel only when sufficient room is available. Among Fibre Channel vendors, both Ancor and Canstar say they support intermix mode in their switches.

Missing Links

If Fibre Channel is to find favour as a network solution, it will need to ensure interoperability among different vendors' gear. Unfortunately, the standard stops short of providing specific technical details about how its features should be implemented. Consequently, "the standard does not provide enough detail for vendors to produce interoperable products," says Ed Frymoyer, Program Manager at Hewlett-Packard.

To address the interoperability issue, Fibre Channel's three most influential supporters — Hewlett-Packard, IBM, and Sun — have formed the Fibre Channel systems Initiative (FCSI). The rather unholy triumvirate aims to develop a set of documents, called 'profiles,' that give vendors implementation information to develop interoperable products.

Three profiles are now under development. The first, the storage profile, will provide design guidance for products used in point-to-point data storage applications based on SCSI or IPI, such as backing up data to disk arrays.

The second profile, the networking profile, will define how adaptors should encapsulate IP packets for transmission over Fibre Channel networks, and also will cover the design of Fibre Channel switches.

To enable Fibre Channel products to interoperate with users' existing LANs, the FSCI's third profile, the internetworking profile, will define interfaces between Fibre Channel equipment and Ethernet, FDDI, and Token Ring networks. "It's important that we are able to support legacy nets," says Hewlett-Packard's Frymoyer. He adds that the internetworking profile also will define interfaces to both LAN and WAN versions of ATM equipment. Frymoyer says the first and second profiles will appear by the fourth quarter; and the third will be completed sometime in 1994.

The fact that none of the FCSI's profiles has yet been published has not stopped vendors (including Ancor, the only vendor now shipping Fibre Channel products) from claiming that their products will conform to the specifications once they arrive. Clearly, a question mark will remain over such assertions until interoperability testing can be undertaken.

Interoperability testing also will be a prerequisite for products handling protocols that are not defined by the profiles, but which are permitted by the options built into the overall Fibre Channel standard.

Fibre Channel vs ATM

ATM now poses an even bigger potential challenge than interoperability to Fibre Channel's widespread acceptance. Nonetheless, Fibre Channel's proponents point to several clear advantages the technology has over asynchronous transfer mode. One obvious advantage is raw speed. Early Fibre Channel products operate at 266Mbps, and 1Gbps products are expected in the fourth quarter; LAN implementations of ATM, in contrast, currently operate at 155Mbps. "The first 622Mbps ATM products are expected in 1994," says Rob Newman, ATM Product Line Manager at SynOptics Communications.

But line speed is only half of the story when comparing network technologies; it's also crucial to gauge latency, or network delay. Fibre Channel keeps network delays lower than ATM in two ways: by doing most processing in hardware, as mentioned earlier, and by using small headers relative to overall frame size.

The maximum frame size in Fibre Channel is 2 Kbytes, including a 24 to 28 byte header, for an overhead of less than 1.5%. ATM cells, in contrast, are 53 bytes long and incorporate a 5 byte header — an overhead of about 10%.

Of course, when working out overhead users must also take into account the size of the packets that they intend to encapsulate using Fibre Channel. Fibre Channel is most efficient when it carries packets that are equivalent in size to Fibre Channel frames, such as those found in SCSI applications. Conversely, using one Fibre Channel frame per small packet "would produce a very large overhead" compared with ATM, says Hewlett-Packard's Frymoyer.

Still, Fibre Channel vendors maintain the problem of small packets can be circumvented by the technique of encapsulating multiple short packets in a single frame containing one Fibre Channel header. However, this solution is only practical for protracted transmissions; it does not improve performance for highly interactive traffic, or acknowledgments.

Another advantage of Fibre Channel's design relates to its use of a hardware-based flow control mechanism to minimise congestion. "This means a frame will never be sent unless there is sufficient bandwidth for it to reach its destination," says Canstar's Malavalli. In contrast, ATM products use a software-based congestion control proce-

SNA OVER LAN – MAKING THE MOST OF EXISTING RESOURCES

RAD Offers the Key to Successful SNA/LAN Integration: Freedom of Choice

With RAD's new family of standalone SNA/LAN converters, you can now migrate your SNA/SDLC devices to multiprotocol LAN backbones without being locked into one particular internetworking strategy.

Supporting Token Ring and Ethernet, RAD's products enable seamless SNA/LAN transport by converting data from cluster controllers, FEPs and other SDLC devices into Logical Link Control (LLC 2) packets and vice versa. You can integrate your SNA traffic over a faster, more robust LAN/WAN backbone.

Consider the benefits of using RAD's standalone converters:

- *Eliminate dedicated SDLC lines*
- *Save on upgrade of IBM-3174*
- *Connect IBM-3274 to LAN*
- *Substitute cost-effective converters for pairs of bridges*
- *Local SDLC polling saves WAN bandwidth*
- *Route SNA/SDLC traffic over IP backbone using any major vendor's router*
- *Save on management costs (and NCP fees)*
- *Upgrade SDLC link speed*

Why be tied to a single vendor solution when RAD offers you the freedom to choose the most cost-effective migration strategy?

SNA/LAN Product Family

STC	*Controller to Token Ring converter, supports up to 2 lines and 15 controllers*
FTC	*FEP to Token Ring converter, supports up to 15 PU type 2.0 LAN connected devices*
SEC	*Controller to Ethernet converter*
FEC	*FEP to Ethernet converter*
RR-STC	*RADring hub module SDLC to Token Ring converter. Full hub supports 10 STC-2 modules, linking up to 150 PU type 2.0 devices*

RAD

data communications

Intl. Headquarters
Rad Data Communications
8 Hanechoshet Street
Tel Aviv 69710, Israel
Tel: 972-3-6458181
Fax: 972-3-498250

Australia
Dataplex Pty Ltd.
P.O. Box 526
234 Whitehorse Rd.
Blackburn 3130, Victoria
Tel: (03)210-3333
Fax: (03)210-3399

Dataplex (Sydney)
Level 6, Eastgardens Office Tower
Bunnerong Rd. & Wenttworth Ave.
Pagewood, N.S.W. 2035
Tel: (02)349 3911
Fax: (02)349 7747

New Zealand
Dataplex (N.Z.) Ltd.
Tel: (09)307 8139
Fax: (09)377 4743

Argentina (1)3314276, **Australia** (3)2103333, **Austria** 222.601011698, **Belgium** (15)290400, **Brazil** (11)8875544, **Bulgaria** (2)725010, **Canada** 416-477-9977, **Chile** (2)5568390, **China** (1)2028393, (852)8922288,**Colombia** 6110111, **Cyprus** (2)449171, **Czech Republic** 422-366251, **Denmark** (42)919555, **Egypt** (2)735415, **Finland** 0 8036033, **France** (1)47118300, **Germany** (89)3189910, (89)716068 **Greece** (1)6547400, **Hong Kong** 5617700, **Hungary** (1)186-8004, **Iceland** (1)681665, **India** (11)2243092, (812)333415, **Ireland** (1)619066, **Israel** (03)6458181, **Italy** (2)27422565, **Ivory Coast** (225)322751, **Japan** 3.52757910, **Korea** 972-3-6458181, **Macau** 853-336676, **Mexico** 525-6828040, **Netherlands** (10)2620133, (15)609906, **New Zealand** (9)3796882, **Nigeria** (1)967598, **Norway** (2)680650, **Panama** (2)13835, **Papua New Guinea** (675)252555, **Paraguay** (21)490076, **Peru** (14)333410, **Philippines** 2.8125068, **Poland** (26)280627, (58)210818, **Portugal** (1)8519980, **Russia** (95)2054749, (95)4375298, **Singapore** 4792088, **Slovenia** (61)122281, **South Africa** (11)8862200, **Spain** (1)6630100, **Sri Lanka** (1)573692, **Sweden** (8)6260650, **Switzerland** (14)917744, (22)3431150, **Taiwan** (2)7716333, **Thailand** (2)2332261, **Turkey** (1)2746894, **U.K.** 908 262121, **Uruguay** (2)923592, **Venezuela** (2)7624624.

dure which results in non-essential cells marked with a cell loss priority (CLP) bit being discarded when the network becomes overloaded — an inherently less reliable mechanism, Fibre Channel vendors say.

But ATM does have some inherent advantages over Fibre Channel. For instance, ATM is able to switch time-sensitive, or isochronous, traffic such as multimedia, video, and voice. Fibre Channel networks, in contrast, can only handle such traffic by using Class 1 dedicated links to ensure that such traffic would arrive on time, and in the correct order.

The still-emerging Fibre Channel standard does not as yet define a way for time-sensitive information sent over Class 2 and Class 3 switched services to arrive in the correct sequence.

"No one knows how Fibre Channel switches will cope with this stuff," maintains Bob Amy, Manager of emerging technologies at IBM.

The X3T9.3 committee already is working on adding isochronous capabilities to Fibre Channel in an addendum known as Fibre Channel Enhanced Physical (FC-EP). "Fibre Channel Enhanced Physical will be part of the standard by 1994," according to Canstar's Malavalli. Isochronous transmissions will be implemented via a new service class called Class 4, he adds.

Another advantage of ATM is that it will be implemented in both LAN and WAN products. When the ATM market picks up momentum, users will be able to build streamlined enterprise networks based on a single technology, without having to convert between frame or cell formats when passing traffic from one domain to another.

Fibre Channel is unsuitable as a wide area transport, since the exchange of frames involved in setting up Class 1 connections delays transmission for an unacceptable length of time.

"The best combination of technologies would be Fibre Channel in the LAN and ATM in the WAN," says Paul Rupert, Manager of advanced telecommunications programs at the Lawrence Livermore National Laboratory, a research facility run by the University of California for the US Department of Energy. Livermore is beta-testing Ancor's VME adaptor and switch.

High Speed, Low Cost

Fibre Channel's success also will depend on its cost. Ancor says its 133Mbps CXT 250 system product costs around $US1,500 to $US2,000 per node, including switch and adaptors. A 266Mbps version costs from $US3,000 to $US3,500 per node. A network that uses Canstar's 266Mbps switch and Emulex adaptors (assuming they will be able to interoperate) would cost between $US2,300 and $US3,000 per node.

The low prices for today's Fibre Channel products are partly attributable to a low-cost fibre optic component developed by IBM. Called the optical link card (OLC), the physical layer device drives the data signal over the fibre line. The OLC has been licensed to Hewlett-Packard and also is being incorporated into products from Ancor, Emulex and Canstar. "The OLC is less expensive than FDDI optical components because the Fibre Channel physical spec is far less demanding," says Emulex's Wayne Rickard.

Relatively low prices are likely to give Fibre Channel a boost in storage and high-end workstation applications. However, "per node cost needs to come down to $US500 before Fibre Channel becomes viable" for LAN usage, says Edward Turner, Director of peripherals product marketing at Sun. Fibre Channel prices are expected to fall dramatically over the next few years as mass production gets under way — but the same is true for ATM equipment.

Nevertheless, Fibre Channel vendors say prices of their products will fall further and faster than ATM, since Fibre Channel is a simpler technology to implement.

"ATM adaptors will probably require between 400 and 500% more circuitry than Fibre Channel adaptors," says Terry Anderson, Senior Vice President of system development at Ancor.

Stephen Saunders

A Router With the Lot

As well as supporting ISDN, frame relay and X.25, the latest ACE multiprotocol router can act as a terminal server.

Interlink Communications is an Australian company which was founded in 1984, about the same time as most of its overseas competitors. The company imports, distributes and manufactures a large range of products based on CCITT, OSI and IEEE standards. The company is particularly strong in X.25 solutions, voice/data integration products and in frame relay expertise. It also provides a range of pre- and post-sales services and counts several large Australian and international organisations among its client base.

Amongst the company's product line is the ACE Router range from New Zealand-based Network Dynamics. In New Zealand, Network Dynamics offers a similar range of products and services to Interlink. Network Dynamics' new design, the ACE Router 2WE-4PM, is being manufactured in Australia by Interlink.

With the addition of the 2WE-4PM, the ACE Router family consists of three basic units. At the low end of the range is the 8WE-4, which provides a single thick Ethernet port with a standard AUI interface, two async ports and 4, 6 or 8 sync ports for WAN interface. Next comes the 2WE-24P, which provides one Ethernet port, two WAN ports and either 12 or 24 async ports, allowing the unit to be used concurrently as a router and terminal server.

Both units' WAN ports can be driven at up to 2Mbps. At the high end of the range, the 2WE-4PM differs from its predecessors in that it is a fully modular unit. This unit's base configuration provides a single Ethernet (AUI) port, four async ports, two sync ports which can be driven at up to 10Mbps, and a single expansion slot. This slot can be filled with one of three option boards, which provide a further 16 async ports, four sync ports plus one Ethernet port or two sync ports, one Ethernet port and an integrated ISDN TA. Both basic and primary rate ISDN are supported. ISDN calls can be initiated on the basis of packet type, protocol type, link quality (e.g. congestion or failure) or a variety of timers such as time of day, week or month.

The base price of all units in the ACE family includes all configuration software. This software supports concurrent routing of TCP/IP, IPX/SPX, X.25 and DECnet traffic. Bridging is fully supported, providing WAN connectivity for NetBIOS-based protocols, LAT, AppleTalk, etc. The units provide full support for PPP (with Jacobson compression), X.25 and frame relay. The frame relay implementation conforms to both CCITT and ANSI standards including extensions such as LMI.

Performance Parameters

The ACE routers are designed to meet the requirements of the bulk of corporate customers and perform accordingly. The 2WE-4PM, for example, can filter Ethernet-to-Ethernet traffic at 14,880 packets per second (using 64-byte frames) over two streams (i.e. using four ports) concurrently. It is also capable of forwarding to its wide area port at between 10,000 and 12,000 packets per second (pps), again using 64-byte frames under test conditions.

While several of the ACE Router's overseas competitors can quote higher performance figures, Interlink and Network Dynamics can point to empirical evidence in Australia and New Zealand which indicates that the units can more than hold their own against such competition.

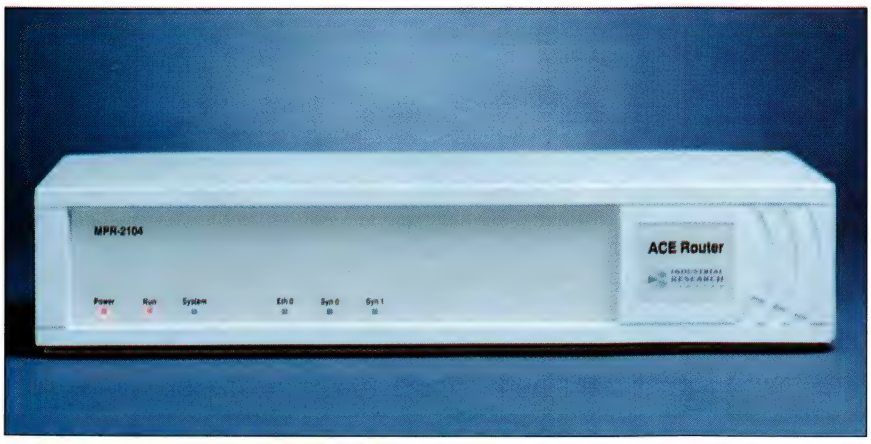

PRODUCT SUMMARY

Name: ACE Router

Description: Wide area multiprotocol router capable of acting as a terminal server

Price: (ex tax) 8WE-4 $15,000 to $19,000 depending on number of WAN ports; 2WE-24P $15,400 to $17,000 depending on number of Async ports; 2WE-4PM $12,000 for base unit. All software is included in base price of units. 2WE-4PM Options: 16 Port Terminal Server Board $4,000; 4 Port Sync Board $5,200; 4 Port LAN [(AUI) IEEE 802.3] Board $5,200; ISDN Board (1 x ISDN TA, 1 x 802.3 and 2 x Sync ports) $10,500

Vendor: Interlink Communications, 4/252 Allambie Road, Allambie Heights NSW 2100 Tel: (02) 975-2577

The reason for this is that filtering and forwarding figures for most routers — including ACE — are derived from tests which, basically, involve attaching a frame generator to the unit under test and flooding its ports with 64-byte frames. Ten minutes with a protocol analyser on a real network will convince anyone that this test procedure bears only a passing relationship to reality. I've seen situations where a router's actual forwarding rate was less than half its manufacturer's quoted rate — the result of a particularly nasty mix of very small and very large packets.

Another point to note when considering a router's performance is that most networks never use anywhere near a router's (or bridge's) maximum bandwidth. Most Ethernet networks in Australia run with an average bandwidth utilisation of 10 to 15%. If you assume an average frame size of 512 bytes at 10% utilisation and (simplistically) apply this to the ACE 2WE-4PM's filtering performance under test conditions (1 x 512 = 8 x 64-byte packets) you will see that the router would be called upon to filter the equivalent of 11,904pps. This is only 80% of its test capability. Similar maths can be applied to the unit's WAN performance. The main point is that there is a significant installed base of the older ACE units in service in real corporate environments, and their owners are satisfied with the unit's performance. Since the 2WE-4PM is a significant improvement on the earlier models it would seem to have a ready market.

Graeme Le Roux

Making Light Work of Frame Relay

The DST 2000 analyses frame relay services to 2Mbps — without difficult decodes.

Most manufacturers would be less than overjoyed to hear their product described as 'lightweight,' but to GN Navtel it's a compliment. To start with, the company's Data Services Tester 2000 (DST 2000) delivers many of the functions of protocol analysers used for frame relay networks — without the burdensome decodes that can leave net managers wishing they'd opted for a graduate degree in advanced electronic theory.

What's more, the DST 2000 tips the scale at under 6 kilograms, roughly half the weight of heftier protocol analysers. And its price tag means fat wallets aren't a prerequisite to purchase. At first glance, the DST 2000 looks — and acts — like a protocol analyser. It attaches to a T1 or E1 line as a frame relay device, filters packets, and summarises the results as statistical analyses of network performance.

But the DST 2000's findings are displayed as easy-to-read lists of, among other things, total frames, short and long frames, number and types of errors, and aborted frames. This easy, straightforward approach should go a long way toward easing the load on anyone charged with installing, maintaining, and troubleshooting frame relay networks.

The GN Navtel analyser also offers a full complement of physical layer tests for leased lines, and it can perform bit error rate tests (BERTs) on T1 and E1 connections or individual channels within them.

Among the likely competition for the DST 2000 is the Digitech Series 900 WAN/LAN Protocol Analyser from Digitech Industries. But the Digitech box can look inside frame relay packets to decode the encapsulated upper-layer LAN protocols. The DST 2000 doesn't offer any packet decoding, but GN Navtel claims that the statistics its unit supplies should be more than adequate for installers and troubleshooters.

The DST 2000 also competes with high-end test equipment from vendors like Telecommunications Techniques or Hewlett-Packard. GN Navtel says that is product's statistical analysis of frame relay exceeds the capabilities of E1/T1 testers.

The Layered Look

The DST 2000 has the ability to analyse the lowest three layers of the OSI stack — physical, data link, and network. Since frame relay packets do not include a network layer, however, the DST 2000 can only offer statistical information about the physical and data link layers.

GN Navtel says it plans to extend the DST 2000 to other high speed digital services such as ISDN PRI (primary-rate interface). At that point, the DST 2000 will be able to take full advantage of its three-layer potential.

The company says the DST 2000 is well suited to field service personnel who work for carriers and service providers, but it's also designed for end-users. "There's clearly a benefit to users having the same test tools as their carrier," says Laszlo Sabo, Senior Product Manager.

The DST 2000 attaches to a network as a frame relay terminal. It can be inserted between the customer network and the wide area frame relay service. The unit features a T1 or E1/G.703 four-wire interface, as well as EIA-232, V.24, V.35, EIA-449, and V.11 data ports.

Once installed, the DST 2000 evaluates service quality by measuring throughput and filtering packets on the DTE (data terminal equipment) and DCE (data circuit-terminating equipment) sides of the network. It counts frame relay virtual connections and lists them according to their status as active, inactive, or deleted.

The unit also filters and analyses packets to report on network traffic. For example, it lists the number of frames with forward or backward explicit congestion notification (FECN or BECN) and summarises information about the frame relay service, displaying the overall percentage utilisation for DTE and DCE and average number of frames transferred per second.

Alert messages are displayed on the unit's screen if the unit encounters invalid messages or inactive or invalid DLCIs (data link connection identifiers).

The DST 2000 also tests the quality of T1, E1, or fractional T1/E1 connections, performing BERTs and loopback tests on all circuits and analysing circuits in detail. It also has a microphone for testing voice channel quality. The DST 2000 can be set to function unattended for 24 hours. All data is displayed on the DST 2000's 40-by-20 character LCD; there also are roughly 60 LEDs on the unit. The DST 2000 is only 15 centimetres high, 32 centimetres wide, and 22 centimetres deep.

Mary Jander

PRODUCT SUMMARY

Name: Data Services Tester 2000 (DST 2000)

Description: A test instrument for installing, maintaining, and troubleshooting frame relay services

Price: From $19,000

Vendor: GN Elmi, 9 Prospect Street, Box Hill, VIC 3128. Tel: (03) 890 6677

Detail Without Decodes

GN Navtel's DST 2000 can be set to perform statistical analyses within the customer premises or on the carrier network. It delivers its findings in easy to read lists rather than the arcane decodes typically used by protocol analysers.

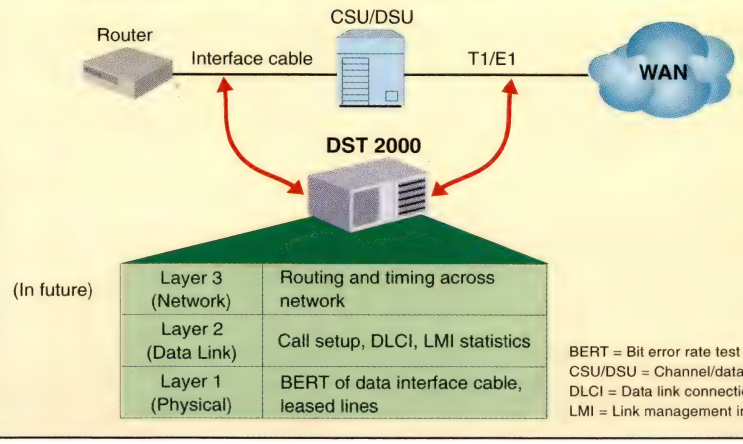

Sending Data Over FDDI Networks

Alfa's FDDI-Sync products guarantee bandwidth for time-sensitive applications.

FDDI, with its bountiful bandwidth, might seem ideal for multimedia and videoconferencing, but bandwidth is only half the battle: these applications are time-sensitive, so they also require guaranteed maximum transmission delays. In other words, these applications need synchronous transmission — a capability supplied by Alfa with its FDDI-Sync products.

Alfa's FDDI-Sync line of adaptor cards and concentrators adds synchronous support to existing asynchronous FDDI networks. The cards can be used with both fibre optic and copper STP (shielded twisted-pair) cabling.

FDDI standards do in fact define a way for the MAC (media access control) layer to handle both synchronous and asynchronous frames, so existing FDDI chip sets theoretically could support both kinds of traffic. But FDDI standards do not yet define a procedure for allocating synchronous bandwidth to stations on an FDDI network. Until now, no other FDDI vendor has added synchronous facilities to its FDDI adaptors.

The American National Standards Institute (ANSI) is working on adding such an allocation feature, which is expected as part of its new station management (SMT) FDDI standard. A draft version could be ready by year's end, after which other FDDI vendors are expected to add synchronous capabilities to their adaptor cards via upgraded software drivers.

Alfa says its new product line will accept software upgrades to conform to the eventual ANSI standard, but it declines to say whether it would charge its users for the change. In the meantime, Alfa hopes to steal a march on its competitors with its proprietary implementation.

The FDDI-Sync line includes adaptor cards for ISA (Industry Standard Architecture) and EISA (Extended ISA) PCs and Sbus workstations, all available in fibre and STP versions. Alfa also sells eight-port concentrators for fibre and copper wiring.

Priority Service

Alfa's FDDI-Sync products ensure that time-sensitive frames arrive punctually by allowing part of FDDI's 100Mbps bandwidth to be dedicated to serving each synchronous workstation on the network. Asynchronous workstations then contend for the remaining bandwidth.

To fully ensure synchronous operation, the cards employ a variation of the token-passing access mechanism used in asynchronous FDDI. In conventional asynchronous FDDI networks, a token is sent around the FDDI ring, stopping at each station to give it an opportunity to transmit frames. The longer a station holds onto the token, the more bandwidth it uses. Timers built in to FDDI adaptor cards measure how quickly the token passes around the ring. If the rotation speed falls below a specified level, the stations assume the network is being run to capacity and stop transmitting.

With synchronous FDDI, the timers in asynchronous stations are configured to assume that the bandwidth dedicated to synchronous stations has already been used up. As a result, synchronous stations will always be able to transmit data, up to their prespecified bandwidth limit. If a synchronous station does not need all of its reserved bandwidth, it simply returns the token to the network early, thus making the remainder of its dedicated bandwidth available to asynchronous stations for the rest of that token circuit.

Synchronous FDDI does pose a couple of tradeoffs: obviously, the performance of asynchronous applications can suffer, and network managers have to be careful not to reserve too much bandwidth for synchronous applications.

Because Alfa's FDDI-Sync products obtain bandwidth using a variation of FDDI's token-passing facilities, they operate transparently to asynchronous FDDI adaptor cards. The new cards will interoperate with third-party vendors' asynchronous adaptors and concentrators, according to Alfa.

Synchronous FDDI products eventually will face competition from a new version of the FDDI standard: FDDI 2, which is being designed from scratch to carry time-sensitive applications such as multimedia and videoconferencing. Synchronous FDDI may also be challenged by ATM (asynchronous transfer mode) as more ATM products appear. However, affordable implementations of FDDI 2 and ATM won't arrive until 1994 at the earliest. In the meantime, synchronous FDDI could serve as an interim solution.

Stephen Saunders

PRODUCT SUMMARY

Name: FDDI-Sync

Description: Synchronous FDDI adaptor cards for use with fibre optic and shielded twisted-pair cabling

Price: (inc tax) FDDI-Sync ISA fibre and STP adaptors cost $3,112 and $2,143, respectively; EISA fibre and copper adaptors, $3,699 and $2,722, respectively; Sbus fibre and copper adaptors, $3,699 and $2,722, respectively. Fibre and copper concentrators cost $14,913 and $9,017, respectively

Vendor: Alfa Incorporated

Distributor: ADE Network Technology, 2 Dublin Street, East Oakleigh, VIC 3166. Tel: (03) 543 2677

FDDI Gets in Sync

Alfa's FDDI-Sync adaptor cards ensure that FDDI frames from time-sensitive applications arrive on time. In this example, two sync stations on a 100Mbps FDDI ring are each allocated 10Mbps. Async stations share the remaining 80Mbps.

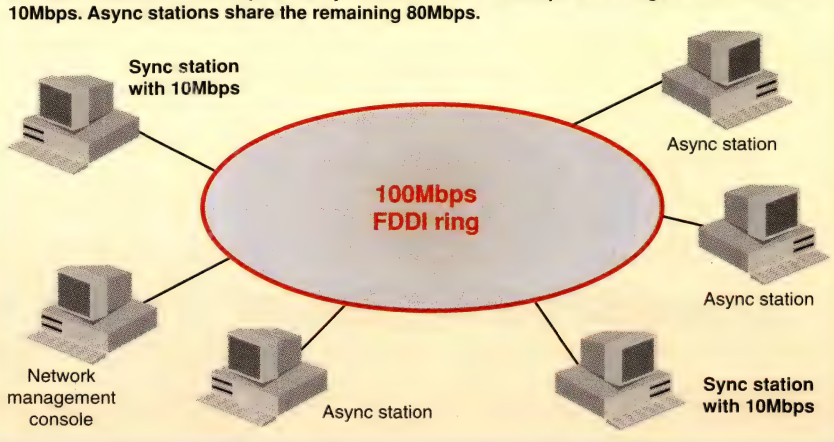

THE ESSENTIAL EDI

An integrated, self-paced learning program meeting the educational requirements of any organization with a need to understand the technical and business implications of EDI.

EDICA
EDI Council of Australia

STANDARDS AUSTRALIA

The Essential EDI is a new, low-cost, high-value learning program specifically designed for managers who want to understand how Electronic Data Interchange (EDI) — the electronic exchange of information between trading partners — can benefit their business, and how to go about the planning and implementation of an EDI strategy.

EDI is a business tool which improves productivity by effectively eliminating delays in document transit, the need to re-key data several times, and the probability of data entry errors. It promotes business-to-business communication, reduces demand on capital by better stock management, and eliminates many other inefficient business practices implicit in a paper-based system.

The Essential EDI contains:

— an introductory video

— a video user's guide

— a glossary of EDI terms

— a 300 page modular workbook covering:
• The business benefits of EDI • Technical options • Hardware/software requirements • Approaches to EDI implementation • Security and legal issues • Essential steps in EDI planning • The value of EDI in changing work practices • The operational benefits of EDI • The financial benefits of EDI

For a detailed brochure about **The Essential EDI**, return the coupon below.

Please send me more information about **The Essential EDI**

Mr/Ms/Mrs _____

Position _____

Organization _____

Address _____

_____ Postcode _____

Send to: Marketing Services, Standards Australia,
PO Box 1055, Strathfield NSW 2135
Phone (02) 746 4700 Fax (02) 746 8450.

EEDIAC 0693

Telecom 1, Optus 0: The 0015 Arbitration

Austel recently decided that 0015 would not be preselectable. Peter Waters reviews Austel's arbitration and looks at the implications of the decision for the future of preselection.

OTC opened the 0015 code in 1988 to provide an alternative pathway to the 0011 code for international facsimiles. Today, about half of all international facsimiles are sent using the 0015 code and the international facsimile market is experiencing rapid growth.

The total traffic volume (both 0011 and 0015) has now grown from approximately 35 million paid minutes in 1987/88 to 90 million paid minutes in 1992/93, and will grow further to a projected 170 million paid minutes in 1997/98 — a 400% increase over 10 years.

'Dial 1' is an interim arrangement which will be replaced by preselection, which provides for the customer's carrier of choice to be programmed into the local exchange so that all long distance and international calls are routed automatically to that carrier without the need for a carrier-specific code. Calls dialled using the 0011 code, including international facsimiles, will be subject to preselection.

Optus argued that as 0011 and 0015 were interchangeable, both should be preselectable. Telstra countered that 0015 was properly regarded as an OTC/Telstra-branded product, and that Optus should establish a new code for its competing facsimile service (like 0018 or 0019).

Government Policy

The Government's basic policy premise is that interconnection should be built around the equal access concept. Equal access is said to ensure that: 'Both carriers have the same opportunity to offer services to all customers and the same access to information systems . . . [and] . . . should also encourage the second carrier to use Telecom/OTC's facilities rather than use its own where doing so minimises the overall cost of providing telecommunications services.' The Minister has stated that "only preselection with appropriate override codes achieves equal access in terms of the Government's decision."

Beyond these general principles, the Government's policy statements are unhelpfully short on the details. The *Telecommunications Act* does not refer at all to preselection, although one of the objects of the interconnect provisions is to remove 'obstacles to consumers having equal access to the telecommunications services supplied by the various carriers.' The carrier licences establish a timetable for preselection to be made available progressively over the next five years, although no guidance is given on what preselection is to look like. The fundamental issue at the heart of the 0015 dispute was which services preselection should apply to. It is clear that preselection was not intended to encompass all telecommunications services and that some would more appropriately be provided behind carrier-specific dial codes, such as enhanced data or other high value added services. The difficulty faced by the carriers, and ultimately by Austel, was to find a basis or 'touchstone' for resolving the preselection status of the range of other services where the issues were less clear cut.

Optus' Case

Optus has argued that including 0011 within preselection while excluding 0015 would cause considerable customer confusion and inconvenience. Many individual customers use 0011 and 0015 interchangeably for faxes, dialling one code if the call does not get through on the other code. Within organisations, some employees may habitually be 0011 users, while others are 0015 users. If 0015 was not preselected together with 0011, customers confusingly would continue to receive accounts from both carriers for long distance and international calls. By contrast, preselection provided a large user with a certain, effective and cheap means of making and enforcing a centralised decision to shift all of its PSTN traffic from one carrier to the other, and back again if there was a competitive advantage .

Optus also argued that the attractiveness of selecting Optus as the preferred carrier, particularly in the upcoming ballot, would be impacted adversely if Optus was unable to offer a comprehensive portfolio of services. While customers who preselected Optus could continue to send international facsimiles by the 0011 code, many customers clearly regarded the 0015 code as providing a technically superior pathway, and the absence of a separate Optus facsimile pathway would persuade customers that Optus was offering an inferior service. Opening and establishing in the market a new Optus international facsimile code by the time balloting occurs in Sydney, Melbourne and Canberra was, in Optus' view, unrealistic.

Optus also rejected Telstra's arguments that 0015 had become associated with the Telstra product 'Faxlink' in customers' minds.

Optus argued that in a monopoly environment, any association between a particular dial code and the monopoly carrier's product accessed through that code is attributable to the carrier being the sole provider. Finally, Optus asserted that exclusion of 0015 from preselection was not consistent with the basic concept of equal access. In Optus' view, equal access was not only about carriers having equivalent access to customers, but also of customers having an equal opportunity for access to the carriers' services, undistorted by factors which have little bearing on the merits of each service, such as inherited dialling patterns like 0015.

Telstra's Case

Telstra argued that equal access is not 'an infinitely elastic concept designed to offset any commercial disadvantage Optus may perceive itself to be experiencing because it is putting its competitive products on the market at a later date [than Telstra].' Telstra recognised that a barrier to competition could arise from the effect of inertia or established dialling patterns when customers choose a telephone service. However, reliance on preselection to overcome these barriers could not be absolute, as preselection imposes significant costs in using network facilities and administering choice.

More fundamentally, Telstra considered that preselection can erode customer choice and blunt the incentives for carriers to differentiate and upgrade their services. In Telstra's view, use of a preselected, or generic, dial code inevitably confines a carrier to a basic 'vanilla' service common to the participating carriers. Separate dial codes provide a basis for and incentive to achieve greater innovation between the carriers. Telstra, in a confidential annexure to its Submission, provided Austel with details of proposed enhancements to Faxlink which it considered under threat if 0015 became a common code.

The basic test, in Telstra's view, was whether the association of 0015 and Faxlink is so exclusive and generic that it precludes Optus from entering the market for facsimile services using another code. In Telstra's view, the barriers to Optus' entry were low. Over half all facsimile traffic was carried on 0011, which would be subject to preselection. Most 0015 users are large companies which can be targeted readily by Optus. Telstra argued that it had put considerable effort into promoting 0015 in the marketplace, and if 0015 became preselectable, this goodwill would be transferred or dissipated and Optus would get a 'free-ride' on Telstra's marketing efforts.

Austel's Decision

Austel's vision of equal access was closer to Telstra's position than to Optus'. Equal access was seen not as an end in itself, but as a vehicle for achieving the principal statutory objective — the promotion of the long-term interests of consumers. While equal access is important in securing for consumers the benefits of a vigorously competitive market, other considerations relevant to the long-term interests of consumers included the continued development and enhancement of the telecommunications infrastructure and the 'user friendliness' of customer access arrangements.

Austel regarded 'equal access' as being more than a term of art principally meaning preselection. Equal access more generally meant an 'equality of opportunity to supply services to all customers.' Austel noted that:

'Equal access is a positive, not a negative concept. It is about 'giving,' not about handicapping or 'taking' something away in order to achieve absolute equality. It is not a theoretical concept concerned with compensating for all the advantages incumbent carriers may have in the marketplace by virtue of their former monopoly or duopoly status. It is a pragmatic legislative safeguard consisting of specific structural measures to ensure that new carriers have the opportunity to compete on an equal basis in supplying services to all customers and to ensure that customers have equal access to the services of all carriers.'

One of Austel's most significant concerns was a perceived adverse impact on customer choice from 'bundling' different types of services together in a single preselection basket. Austel stated that:

'If significantly different services were included within a common preselection basket and the choice of carrier was dictated primarily by the characteristics of one of the services, a carrier could be acquiring market share in other services without even having to actively market those services. Under such circumstances, a customer's choice of carrier for the other services might have little to do with the price, features or quality of the preselected carrier's version of those services.'

If 0015 and 0011 were both preselectable, customers would not have the opportunity to 'cherry-pick' 0011 and 0015 services between Optus and Telstra.

As there was not the technical capacity within the Telstra network to allow this to happen through preselection, customer choice could only be preserved through establishing a separate dial code for the competing services. Austel then was faced with the choice of leaving Telstra with 0015 or requiring both carriers to obtain a new code, so each was equally disadvantaged. While Austel acknowledged Telstra's advantage of inherited dial codes, Austel stated that an initial competitive advantage due solely to past customer behaviour and not to any future inequality in the carrier's competitive opportunities, may not, in itself, be considered to have an anti-competitive impact on customer choice. Austel accepted Telstra's view that, with the rapidly growing facsimile market and the sophistication of users, any advantage of 0015 would soon be overcome by Optus.

Austel did require that Telstra move quickly to condition its network to establish any separate Optus dial code for facsimile services. Austel also set out guidelines for the future which provided that a particular service should not be included in the preselection regime unless:

- The service is a basic carriage service;
- The service is available to a substantial proportion of the Australian population;
- The service is offered by all, or the majority of, carriers that offer services on a national basis; and
- The carriers' different versions of the service essentially are identical in core functionality, though they may differ in additional features or quality.

In Austel's view, these criteria essentially meant that only services that are well into their product life cycle should be included in the preselection regime unless the carriers can guarantee that over time a new service will meet these criteria.

Conclusion

There are issues and difficulties arising from the guidelines which will need to be worked through. The primary requirement that the service be a mature service to be subject to preselection presents something of a 'Catch 22' situation: by the time a service reaches maturity, the service-specific code may have become identified so closely with a particular carrier that it is not suitable for preselection. If the 0015 code, established for over five years and accounting for nearly half of all facsimile traffic, is not a mature service, then what are the criteria for regarding a service as mature?

What degree of differentiation is permitted between common core functionality and carrier-specific 'add-on' features before the carrier services can be regarded as too different to be preselectable, given that 0015 shares some level of functionality with 0011? However, putting to one side the merits of Austel's decision on the 0015 code itself, the arbitration decision, particularly the annexed guidelines, provides clearer guidance on the meaning and ambit of preselection.

Peter Waters is a partner at Gilbert & Tobin, specialising in telecommunications and technology contracting. The views expressed in this column are not necessarily those of any of the firm's clients. The author was involved in the preparation of the Optus case.

IT'S HARD TO GET A PHONE CONNECTED WITHOUT GOING THROUGH JNA.

Telecom's customer service is really switched on with the help of network communications technology from JNA.

This technology, integrated with Telecom's AUSTPAC service, forms what is arguably Australia's largest OSI network. Operators throughout Australia now field all kinds of customer enquiries from a single desktop PC, accessing and exchanging data from three different computer systems - IBM, Honeywell and Digital.

Previously, three separate terminals would have been needed to do the same job.

JNA's Network Integration knowledge is at the forefront of today's climate of open environment computing. JNA's engineers are the quiet achievers assisting Australian industry for over 30 years.

If your needs are to distribute the computing power of your organisation, a phone call to JNA could have you well connected.

Sydney 417 6177 Melbourne 532 8555 Canberra 248 5422
Brisbane 870 3434 Perth 481 4034 Adelaide 338 2511 Darwin 81 1499

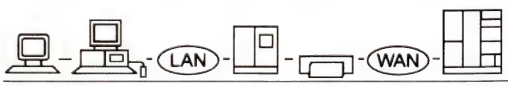

JNA

J401

POINTING THE WAY TO CLEVER COMMUNICATIONS

TOWARDS 2000
TUANZ 93

The Telecommunications Event that will prepare you for the year 2000!

This is the chance for business people to get the information they need in a non-pressure environment.

TUANZ 93 will be the largest telecommunications event ever in New Zealand, featuring the latest technologies from around the world.

Enjoy a variety of presentations, displays, workshops, seminars and social events where you can meet people in this vital industry.

You can't afford to miss

TUANZ 93 CONFERENCE AND EXHIBITION

Ellerslie Function Centre
Auckland, New Zealand
4-7 August 1993

FOR FURTHER INFORMATION CONTACT
Telecommunications Users Association of New Zealand Inc.
PO BOX 331 014, Takapuna, Auckland, New Zealand.
Telephone +64 9-488 1602 Fascimile +64 9-489 9515

Premier Sponsor

THE DOMINION
IT InfoTech WEEKLY

Gold Sponsors

BCL

BELLSOUTH

CLEAR

NEW ZEALAND'S INFORMATION SYSTEMS MANAGEMENT W
COMPUTERWORI

Telecom

TELEVISION NEW ZEALAND

How Britannia Plans to Rule the Telecoms Waves

Peter Hutton
Managing Director, BT Australasia

Peter Hutton is Managing Director of the Australasian subsidiary of British Telecommunications PLC. He joined BT, then the British Post Office, in 1961 and was Deputy General Manager of the Cambridge Telephone area when he left to work for the Hong Kong Government as a telecommunications adviser in 1979. Prior to rejoining BT and moving to Australia in 1987, he was General Manager of the China Telephone Co Ltd, which operated a public mobile radiotelephone service in Hong Kong. He is Chairman of the British Direct Investment Committee for the Australian British Chamber of Commerce. Liz Fell spoke with him last month at BT Australasia's Sydney headquarters.

How do Australia and New Zealand fit into BT's global ambitions?
Hutton: BT is focused on Europe, North America and the Asia-Pacific — the three regions where existing or potential multinational customers are headquartered. Within Asia-Pacific, we are located in Japan, Korea, Hong Kong, Singapore and Australasia.

AC: *Are you targeting mainly British-based multinational companies with offshore subsidiaries?*
Hutton: The first group of customers we are targeting is the top 1,000 or so multinational companies in the world who want first class, end-to-end international services. These companies may have headquarters in the UK, Japan, the US or Australia. So we're not just servicing British companies offshore.

AC: *Do you expect these customers to spend a specified amount per year on telecoms?*
Hutton: We're talking about very big global customers, the IBMs and so forth.

AC: *Are you focusing on particular vertical market segments such as financial services or manufacturing?*
Hutton: We've segmented the market in the normal way. Some sectors tend to be more telecommunications-intensive internationally at this point in time. The finance sector is very intensive, for example, and manufacturing is less so. But as manufacturing changes, so telecommunications needs could change as well. It's wrong to just pick a sector and say that BT wouldn't be interested because it doesn't bring in certain returns. We're looking at a range of sectors and concentrating on some at this stage because they tend to be the ones that will bring us the customer base. It's not purely revenue, it's telecommunications-intensity.

AC: *How will the new BT/MCI joint venture impact on BT Australasia and its customer base?*
Hutton: The new BT/MCI venture is an important strategic investment in North America and gives BT the ability to provide our customers with more products and better service and reach. We have already made good progress in North America, but to realise our ambitions fully we have always recognised that we need a strong American partner, not least to leave us with more freedom to concentrate our resources on Western Europe and the Pacific Rim. MCI fits the bill perfectly.

AC: *During the last nine years, isn't it correct to say that BT has pursued several different offshore strategies?*
Hutton: BT looked at a number of options from the time it was privatised in 1984 until 1987-1988. It entered the PABX manufacturing field by taking an equity stake in Mitel in Canada, and looked at all sorts of activities, from ownership of carrier facilities to product sales. With increasing competition and privatisation BT went through a change of culture which culminated in a company vision, new values and new corporate image.

Within the UK, we have a general carrier licence and mobile operation. Outside of the UK, we concentrate on servicing the needs of multinational customers and we extend our reach and resources by providing services for large domestic customers in some countries.

AC: *Is the NSW Government network contract an example of this?*
Hutton: Obviously this is not an international network, but it gives us infrastructure, not in the purchase or ownership of plant in the ground, but in resources in terms of size of company, the ability to undertake activities professionally and ownership of the access network to international nodes.

AC: *How does BT's search for mobile licences fit into this strategy. You recently sold your stake in McCaw, the US cellular operator?*
Hutton: BT has in the past bid for mobile licence opportunities in a number of countries. This was one way of entering the market.

AC: *And mobile is no longer a BT entry strategy?*
Hutton: It may be an entry profile for some countries, but it is not a major thrust.

AC: *At one stage, BT was following Cable & Wireless around former UK colonies in the South Pacific: Wherever C&W had a franchise, there was BT. Was the aim to destabilise your domestic competitor in offshore markets?*
Hutton: Not at all. If you look at Australia prior to the Optus involvement, C&W was really a small office. They were here bidding for contracts. BT's strategy in Australia was to establish a full-blown office, attract customers and start to build infrastructure. We had two completely different strategies on how to build a customer base.

AC: *Did the 1989 purchase of Tymnet from McDonnell Douglas mark the first step in BT pursuing its global ambitions?*
Hutton: Tymnet was a very significant step for us, the first of our international platforms. It gave us the beginnings of a worldwide data network, which we've since built on, and a customer base, particularly in the US.

AC: *Did it deliver a customer base for BT in Australia?*
Hutton: It wasn't huge. When the Tymnet purchase went through, there was one large X.25 private network and Infolink, which was a domestic public network and a few customers.

AC: *Tymnet became the platform for your Global Network Services portfolio. How big is your GNS customer base here?*
Hutton: We've taken on board a considerable number of customers. In the early days, the major growth came from customers who had been taking the service in North America, for instance, but had a branch office in Australia. That has now been reversed and we're taking on customers here who want service to the rest of the world.

Our biggest single contract is with the New Zealand Ministry of External Relations and Trade, which extends to about 60 countries. And we still have those two original customers. We have an interconnect agreement with Infolink so they can resell BT's GNS to their domestic customers.

AC: *What is a 'considerable number' of customers?*
Hutton: I don't want to give precise figures because we're in very strong competition in this field. But we have a growing customer base.

AC: *Building on the GNS portfolio, is it correct to say that BT has expanded into global voice services, including the provision of international switched voice services?*
Hutton: We started with packet switched data because deregulation normally took place in the data networks in advance of voice networks. Now we're seeing the deregulation of voice, so we put in a global platform for voice/data and even video for Syncordia, which provides outsourcing solutions for the world's largest corporations. Syncordia has access to what I call the Syncordia network, which is basically point-to-point with bandwidth flexibility. Our next network is an international virtual private network to switch voice, data and video which allows for flexibility of bandwidth switching and so forth. This is a complementary network so that we can offer the full range of network services to international customers.

AC: *Is this the international network that was code-named Project Cyclone when it was first leaked in the US media last year?*
Hutton: Yes. It's another global platform which adds to the facilities that BT can offer to customers.

AC: *When Cyclone was originally revealed, switching centres were to be located during 1993 in London, New York, Frankfurt and Sydney. Is the Sydney hub in place yet?*
Hutton: We have an international node in Sydney, but we're not offering customer service yet. At the present time we're still involved in the process of carrying out implementation trials.

AC: *Has this new network acquired another name other than the code-name Project Cyclone?*
Hutton: It's known well as Project Cyclone, although we call it an international virtual private network.

AC: *It's interesting that the name Project Cyclone is now so well-known.*
Hutton: It shows the level of interest in new competitive services throughout the world.

> *"I think that as the telecoms market enters a phase where we have to cooperate with our partners as well as compete, the important thing, and it's something I believe very deeply, is that this be done in a professional manner in terms of market segmentation, and with an understanding that the old monopolies around the world have gone forever."*

AC: *BT seems to be good at creating the type of media hype that has accompanied both Cyclone and Syncordia. Perhaps this is a clever marketing strategy by BT?*
Hutton: Perhaps it's the press that is quite good at getting hold of information.

AC: *Does Cyclone essentially involve international simple resale in that you resell leased private line circuits from Telstra to provide switched services?*
Hutton: One can become an international reseller simply by taking a leased circuit and putting the traffic of more than one customer down it. The Cyclone network is more than that. It is offering a service to customers by giving them the facilities they cannot get through the standard telephony and data networks of the world. International simple resale involves renting a link, dividing it into 20 channels, and selling off the 20 channels for more than was paid for the link. I want to get away from the notion of simple resale because Cyclone, Syncordia and GNS are not simple resale opportunities. They take line plant, connect it up to switches, and then offer a total service to customers.

AC: *What facilities will Cyclone offer?*
Hutton: We'll provide end-to-end connectivity and responsibility, specialist billing, bills in one country and in one currency, certain numbering facilities such as linked numbering plans, flexibility of bandwidth-on-demand etc. There are a number of functions of that nature.

AC: *Would you agree that the removal of resale restrictions on switched voice is the cutting edge of global competition because it offers an inexpensive way for foreign carriers to enter the home markets of their international partners?*
Hutton: We've been competing since we came here in 1987 in one form or another and, at the same time, we cooperate with Telstra as a partner on international carriage. We are an extremely large customer of Telstra in terms of our leased links. I think that as the telecoms market enters a phase where we have to cooperate with our partners as well as compete, the important thing, and it's something I believe very deeply, is that this be done in a professional manner in terms of market segmentation, and with an understanding that the old monopolies around the world have gone forever. I believe market segmentation is inevitable. It will continue around the world until it settles at some level.

It's interesting to compare the telecoms industry with the oil industry. Right from the start, the oil industry, which is a young industry, cooperated in certain areas and competed in others. Looking at it from a marketing point of view, at the crude oil shipment and processing level they cooperate by using each others' tankers and filling up at the same depots. But at the retail end, they fight like hell. They've segmented the market that way. In the telecoms industry, the monopolies have always cooperated fully across markets. We're now entering a phase where we will compete in certain market segments and cooperate in others. You can't send a call from Australia to the UK without cooperating with someone at the distant end. Unless every single PTT sets up its own network in every country, which it won't do, it has to cooperate for those purposes.

AC: *In which market segments will supercarriers such as BT/MCI seek to compete, or is it too early to say?*
Hutton: If you look at the market today — and it's going to change — but if you look at the ordinary IDD voice streams and the standard leased circuits that customers want, then we cooperate with all the PTTs around the world, including Telstra, in providing end-to-end service. But the packet switched networks have never been monopoly networks. The fact that they're growing worldwide is a sign that packet switching is accepted as a competitive service. Now we're seeing competition extended to areas where PTTs are beginning to feel uncomfortable.

AC: *I assume this discomfort is illustrated by the arguments between BT and AT&T over providing switched services via inter-*

national private line resale in each other's markets and, in particular, the transparency of BT's UK interconnection arrangement?
Hutton: In interconnection, BT has an arrangement which is actually overseen by Oftel, the UK regulator. Oftel approves all of the interconnect rates that BT uses and has the right to investigate any problems or disputes. I understand that AT&T recently requested these rates be published, and BT has said that it is happy to do so, whether or not this is a requirement of the UK legislation.

AC: *Now Telstra has entered the fray with the overseas media reporting it may try to block the Australian hub of Cyclone if it doesn't get 'commercially equivalent' access to BT's network. What does this mean?*
Hutton: I don't know what they mean by 'equivalent.' But if Telstra wants interconnection to BT's network, then it is available and it has been given to other telecommunications suppliers over the last two or three years. I don't know what Telstra has asked for in the UK so I can't comment on it.

AC: *Some of the complaints, particularly from US long-distance carriers, suggest that BT is offering access to itself at lower charges. Would you comment on that?*
Hutton: In the UK, BT has been told to eliminate cross-subsidies between services so that competition can be seen as fair and equal. Access charges are controlled by Oftel and BT is bound by regulation not to discriminate.

AC: *Telstra is a shortlisted bidder for the UK Government Telecommunications Network [GTN]. Did BT apply for this contract?*
Hutton: No, BT did not submit a bid.

AC: *Should Telstra win the UK GTN contract, then each of you will be stepping into the other's home market.*
Hutton: That's as it should be.

AC: *I understand that as part of the GTN bid, Telstra has applied for a public telecoms operator [PTO] licence which would be for domestic services, since the UK Government is not awarding international PTO licences. Is that correct?*
Hutton: I'm not aware of this application. But it is worth noting that the legislation in the UK is different, therefore the type of licences issued must, by definition, be different. I think that it will be a long time before we get to a situation where countries can actually look at their legislation and say it is the same and therefore the licences are the same. If you're looking for competitive options in the UK, I would say the UK is more open to competition than Australia at this point in time. In the UK, there is no duopoly any more. If you get a full-blown operator licence, you can put down your own cables, run networks and be a completely separate entity. There are in excess of 30 cable television companies that run their own cables and their own services. Until 1997, a company cannot lay plant here.

AC: *Do you think the Australian Government will open the market in 1997?*
Hutton: That is a matter for the Government. I'll wait and see.

AC: *In relation to the NSW Government's proposed 'corridor' of line links, was BT upset when the Federal Government intervened to protect the carrier duopoly?*
Hutton: No. It was irrelevant to us.

AC: *Under this contract, what is the relationship between BT and the NSW Government Telecommunications Unit [TCU]?*
Hutton: The TCU is the Government's arm for ensuring that the contract is delivered and for looking strategically at the future in terms of what telecoms services might be needed by Government agencies and departments. We see our relationship as a partnership.

AC: *Was BT consulted, for instance, before the TCU signed a Strategic Partnership Agreement (SPA) with Telstra?*
Hutton: We've developed a partnership with the TCU during the long drawn-out tender period, and obviously more so since we've been awarded the contract. We sit with and advise them on what we think their position should be, though they make the final decision. We were part of the negotiations with Telstra on the special arrangement. So we have a very close working relationship which has worked well. If it were purely a contractual relationship of supplier and customer, it could lead to problems.

AC: *After being offered the contract, Horizon is reported to have opted out because it wasn't commercially viable. I assume BT did not agree with this assessment?*
Hutton: I don't know that Horizon withdrew their bid. The Premier said that Horizon had, in fact, been the preferred supplier and had failed to sign the contract. What happened was that the Government went out for 'a best and a final' offer to all three bidders again, and it was on that 'best and final' we actually won the bid.

AC: *Did BT pursue the NSW Government contract for revenue or strategic reasons?*
Hutton: It was both a strategic and revenue-generating opportunity.

AC: *Do you expect to make a profit on the contract?*
Hutton: Yes. We put in a bid which made a return on capital for us as a standalone project. BT does not go in for buying market share. It looks at a project and asks whether it is worthwhile. This one was.

AC: *What do you mean by a standalone project, given BT's global reach and scale economies?*
Hutton: BT viewed the project in the light of its opportunities in Australia. The bid was, however, fully content to provide a reasonable return on capital and was not cross-subsidised from other business.

The NSW Government project gives BT two things. The first is an increase in the reach of BT's infrastructure outside the Sydney metropolitan area, so that we can pick up other customers on that network and bring them in to Sydney from Newcastle, for example.

AC: *So you could transport third party traffic such as the traffic of BHP?*
Hutton: Yes. Third party was part of the contract. The other reason is that we were 60-people strong here and growing at about 20 people a year with normal business growth. That didn't give us the ability, on all occasions, to service our customers.

The NSW Government contract will give us a significant increase in people and resources, with the skill base we need to service other customers. That gave us the reason to put in the major part of BT's networking facility here. We will have a Network Management Centre in Sydney that will look after all of BT's networks throughout the Asia-Pacific.

AC: *Presumably the long history of British investment in Australia makes this an important base for BT?*
Hutton: Once you've decided to become a global player in the three regions, Australia becomes a very important focal point, especially if you look at American and Japanese investment here as well.

AC: *Is BT Australasia bigger than BT in Japan?*
Hutton: Size of the company is not a major indicator. The market in Japan is significantly different from that of Australia and one has to provide an organisation to serve the local customer environment.

AC: *Australia seems such a tiny market that it's difficult to understand why the global telcos are so interested in coming here, though perhaps we could expect BT would be watching C&W closely?*
Hutton: We're not here because of Cable & Wireless. Let me give you some figures. About 70% of BT's largest customers have a presence in Australia, either a branch office, subsidiary, partnership or association. Those companies employ about 175,000 people in Australia. Australia is a significant player in the Asia-Pacific marketplace for BT because of its customer involvement.

Liz Fell is a freelance journalist based in Melbourne.

50bit/s to 2048kbit/s: Hand-held test solutions for Data, PCM, ISDN and GSM Networks

PF-30 Bit Error Analyzer

- 50bit/s to 2048kbit/s
- Interfaces:
  Built-in — G.703 2048/704kbit/s, G.703 co-dir., V.24/RS232C, V.11/X.24
  Via adaptors — V.35, V.36/RS449
- Remote control

PFA-30 Digital Communications Analyzer

- 50bit/s to 2048kbit/s
- Interfaces:
  Built-in — G.703 2048/704kbit/s, G.703 co-dir., V.24/RS232C, V.11/X.24
  Via adaptors — V.35, V.36/RS449
- Framed generator
- Framed receiver
- Mux/demux
- Remote control

PF-25 Bit Error Analyzer

- 1kbit/s to 72kbit/s
- Interfaces:
  Built-in — V.24/RS232C (sync)
  Via convertors — V.11/X.24 (via IC-21), V.35 (via IC-35)

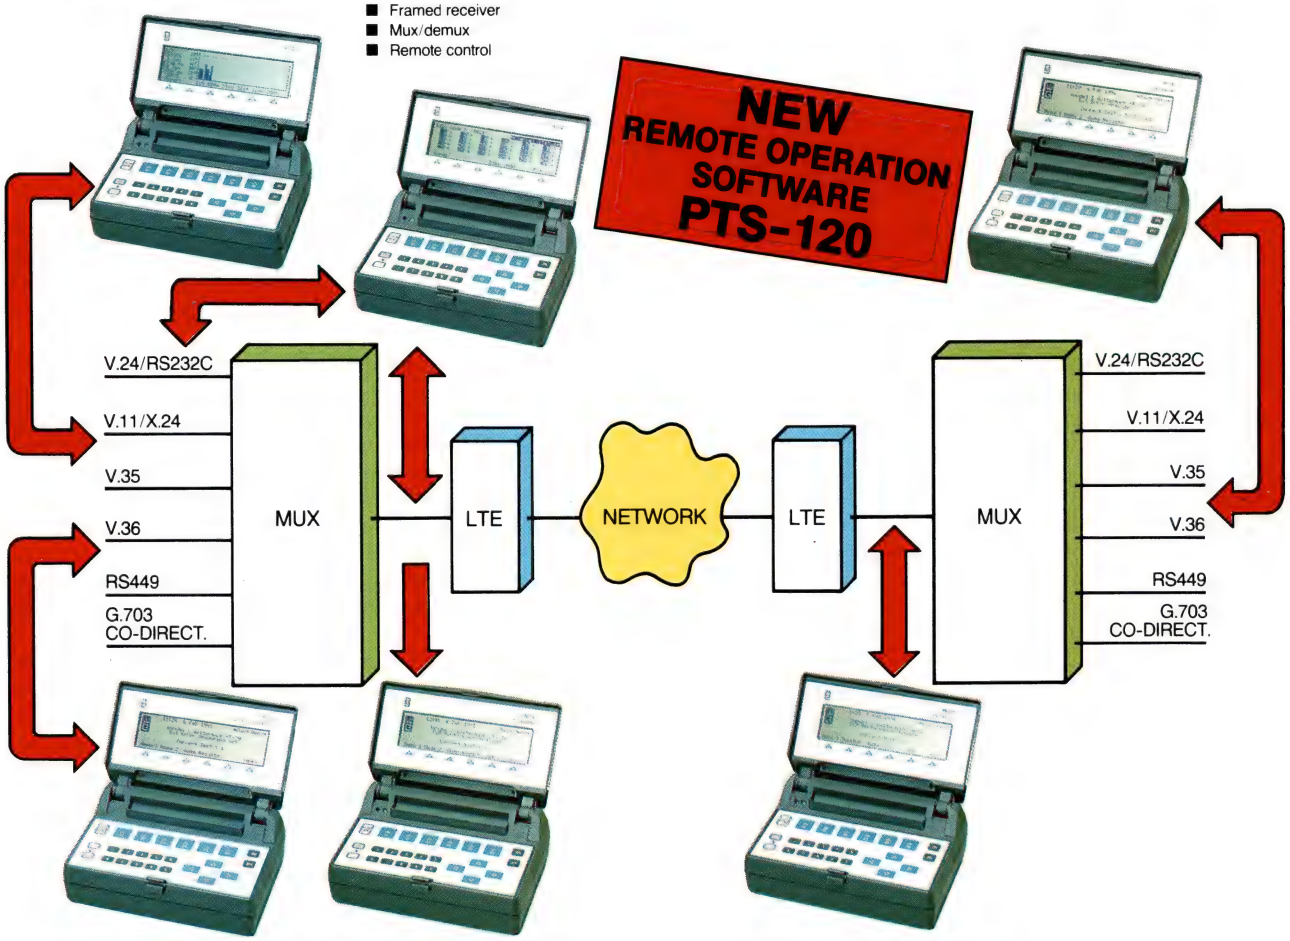

PF-20 Bit Error Measuring Set

- 1kbit/s to 256kbit/s
- Interfaces:
  Built-in — G.703 co-dir., V.11/X.24
  Via adaptors — V.35, V.36/RS449

PA-10 PCM Performance Analyzer

- 2048kbit/s
- Interfaces:
  G.703 balanced and unbalanced
- Unframed generator
- Framed receiver
- Remote control

PA-20 PCM Performance Analyzer

- 2048kbit/s
- Interfaces:
  G.703 balanced and unbalanced
- Framed generator
- Framed receiver
- Remote control

The Wandel & Goltermann family of hand-held BER and PCM testers offers multi-purpose, multi-interface test solutions for commissioning, maintaining and trouble-shooting digital networks.

The user can carry out framed and unframed testing on a wide variety of equipment such as multiplexers, demultiplexers, digital cross-connects, automatic protection switches and circuits operating at speeds from 50bit/s to 2048kbit/s.

Applications include the evaluation of the quality of digital transmission paths, trouble-shooting data circuits and isolating network and transmission faults.

The new Performance Test Software PTS-120 enables the PA-10, PA-20, PF-30 and PFA-30 to be operated remotely.

Wandel & Goltermann Pty Ltd
42 Clarendon Street
South Melbourne Victoria 3205
Tel: (03) 690 6700 Fax: (03) 690 6750

Wandel & Goltermann

Electronic Measurement Technology

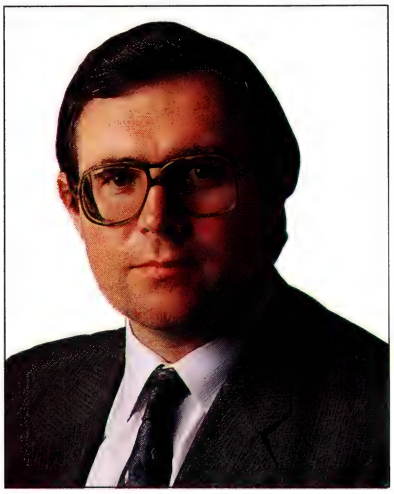

Tom Amos

Charging All Out

Unit fee telephone and data calls across Australia are fast approaching as Telecom moves quickly to re-balance the long haul/local call reverse arrangements in the face of increasing competition.

The mooted introduction of A-party-based radial billing distances and the removal of the current charging boundaries is currently underway in Telecom as a potential competitive response to the limited success of Optus in the long haul and mobile telephone markets. It is difficult to understand how one would define success, but Telecom obviously is not waiting around to discuss the fine details of the definition, and has developed a proactive response that has far-reaching implications for the whole industry.

Currently, local calls within a charging zone or boundary are untimed. This is locked into Government policy, following a number of attempts to introduce timed calls by Telecom over the years. Each attempt was carefully timed to impact the maximum political clout possible, and over the last 10 years Timed Local Calls (TLC) can be credited with a significant voter impact when required.

It was last ruled off the agenda for residential use by the Prime Minister in 1990, although he carefully allowed timed calls to be imposed for business uses on an as-required basis. Further services, such as ISDN Faxstream and so on, are all timed.

Telecommunications carriers have always sought ways to recover the maximum revenue from the user, whether it be for a voice, data or image service. This fiscal recovery has little to do with the cost of provision of the service, and more to do with the perception of what the market will bear. Timed calls are simply another way of extracting a buck in a potentially fairer way, so as to match the revenue needed to provide a service plus a small margin. There must be another way!

It is now proposed to circumvent the TLC embargo by overhauling the billing and charging zone system to represent the time and distance of the call from its point of origin. Charging zones would be a thing of the past, and the tariff plan that applies to each customer would potentially be unique. For example, you could have an agreement with the local carrier for so many kilometres as a unit fee, but pay a higher STD and IDD rate when using the longer distance services. Alternatively, you could elect to go on a fully timed rate and pay lower charges for all services including the STD and IDD components. This would mean tariff balancing like you have never imagined — an Austel nightmare!

'Family and Friends' is just a marketing gimmick (albeit an MCI-based one) to get the public used to being able to select a package of discounts that best suits them. But what of the concept that created the new industry, Basic Carriage and so on? So much for BCS/HLS arguments of the past — now we'll have service blurring that will be almost impossible to unravel. And confusion that will no doubt extend Telecom's consistent rating as Top of the Pops at the Ombudsman's office. If you thought having an after-hours discount was difficult, just wait.

But a discount is a discount, and under the current arrangements it is illegal for the dominant carrier to offer selective discounts on basic services . . . they must be offered by way of an open rate card to all users. Is that simple?

Radial-based charging very smartly gets around the community area calling issue by abolishing the absurd charging zone anomaly, and this is a very good development. What it does also, of course, is neatly change the intent of the Act and the playing field rules that are now in use.

At the same time Telecom is attempting to have the 1991 *Telecommunications Act* amended so to allow the FlexiPlan concept to be retrospectively incorporated — good move! Almost all services, including the product names, are currently BCS; they are basic services, and were ruled so by Austel in those 'blackletter law' readings of the Act last year. How that occurred still mystifies me, but it looks now like Telecom was a little too clever in declaring almost all its services basic BCSs for protection. It may not now be able to meet the competitive challenge with an all-BCS service mix on a head-to-head basis when the competitor is allowed to discount BCSs as it sees fit.

The game rules were well known at the time, and there were several months of backroom manoeuvring prior to Austel accepting the original Telecom position. Changing now because it seems a good idea could hardly be called fair play. But the larger picture is the complete change of charging principles to a new basis.

The opportunity to change the role of new carriers after market entry would be a clear change in both the industry terms and the direction that deregulation is going in Australia — heady stuff! A new billing industry would spring up overnight to validate the new bills, but what would be the cost? All the current billing and TIMS Systems would be redundant, and the arbitrary balance between local and long haul calls would be easily changed, so that the long haul component would only be a small portion of the total calls.

A regime of free local exchange calls could very well work with this type of scenario. The choice of low access fee against a rebated untimed access system would be, of course, an interesting debate, but I bet that the final analysis would show most household users better off. This is an interesting problem in itself, as the area for competition in long haul would in fact shrink markedly, and the US model of a regional monopoly may eventually emerge. Whether there would be any actual reduction in the overall bill to the user is questionable, but it would definitely enshrine this ludicrous concept of two bills, should you choose to use the alternate carrier. This concept is gaining popularity with the carriers. Overhauling the charging systems to reduce the cost of calls to the user is a good thing but is this overhaul for the good of the user, or the good of the carrier?

Tom Amos is a partner with telecommunications consulting engineers Amos Aked Swift.

Kevin Morgan

The Ballot Sham

Austel and the two carriers have presented the Australian telephone user with a fait accompli on the preselection ballot. The 'in club' agreement will see millions of dollars squandered over coming months, with the first round of the charade starting in Canberra in mid-July.

Neither the carriers nor Austel consulted with the consumers who will foot the bill before they locked themselves and the Government into a ballot which threatens to make the telecommunications industry a laughing stock.

No other industry in Australia has ever sought to split the market in such a crude manner, and consumers can only ask how did this latest debacle in the communications portfolio come about.

Austel publicly ducked responsibility for taking the balloting path to preselection in a recent edition of the ABC's *Lateline* program, whilst Telecom argued that it only reluctantly agreed to the process at the regulator's behest. The Government has stood back and said it was merely a matter for the carriers and the regulator to decide upon, leaving Optus as the unashamed public champion of the process.

Optus claims that it had always understood that a ballot was to be part and parcel of the preselection process. This was news to others in the industry, including the CWU, who had expected that preselection would come through a market driven 'canvassing' campaign which would see the market — not regulation — determine the balance between Optus and Telecom.

Consumer groups are still up in arms at a decision which they see as anti-competitive and unnecessary, but in a competitive environment in which the consumer is supposed to be sovereign, no-one is listening to their views.

How then did an industry which prides itself on cleverness get into this mess? It would seem that prime responsibility must fall on Austel, which, despite public coyness, has been the enthusiast behind the balloting process.

Austel's position on balloting would seem to flow from a profound misunderstanding of the relevance of an American-style preselection ballot to the Australian market, which is based on end-to-end duopoly, not discrete market segmentation.

Had either Austel or the carriers stopped to consider the unique circumstances in Australia, then it would have been readily understood that preselection ballots conceived in the post-divestiture American market had absolutely no merit or place in the Australian environment. In the US, structural separation left consumers without a direct relationship with a long distance carrier. Consumers had to 'choose' a long distance carrier either explicitly or initially by default by not returning their ballot paper.

Balloting was seen as a tool to spread market share amongst long distance carriers, and AT&T lost its initial default status and suffered proportional allocation of non-respondents to its competitors when the FCC deemed that insufficient market share was being transferred to competitors.

Despite losing 20% of the market, the process was not all bad news for AT&T because, with preselection, AT&T's competitors were required to pay the same interconnect rate to the local exchange carriers as that paid by the dominant carrier.

Preselection in Australia was also to be the trigger for commercial interconnect rates, but in a staggering strategic blunder Telecom has failed to couple renegotiation of interconnect rates to the balloting process! Telecom is allowing Optus to enter the preselection era with subsidised interconnect rates that will add hundreds of millions of dollars to the losses Telecom will suffer.

Increased interconnect rates are only part of the missing preselection picture in Australia. In the US, consumers also had to pay more after preselection, with the value of falling long distance prices being offset by hikes in local rates through both a Subscriber Line Charge and increased local rates.

Such tariff rebalancing was deemed necessary to maintain affordable universal service and maintain funding of the local loop. Again in Australia the rebalancing element of the preselection equation is missing, and if Telecom loses heavily in the ballot, funding of the local network would be clearly at risk because the STD revenues that subsidised the local network would be lost.

Whether Austel realises it or not, Australia is not the US. In the US consumers had to choose a long distance carrier, but real choice was offered in the market because there were multiple carriers and resellers on the ballot paper. In Australia there will be no multiple choice of carriers. Consumers are confused about the market structure and are unclear about how service will be delivered in the preselection era. The ballot will heighten this confusion.

Nor is there any need to use a ballot to give extra regulatory favours to Optus. The weight of existing competitive safeguards is already such that Optus could hardly fail in its selective market entry. The ballot is a costly sham that consumers should not bother responding to unless they wish to change to Optus as their long distance carrier. This, of course, is the only issue before consumers, because the agreement on default customers remaining with Telecom means that there is no need to register a positive choice for Telecom. The simple fact is that consumers are not being asked to choose between carriers, they are simply being asked to change to Optus. Had Austel given any thought to the real nature of this question then it would quickly have understood that it was a question for Optus to pose directly to consumers and not a question that had to be put through a ballot.

But with Austel fixed on the goal of promoting competition at any price, such obvious facts are overlooked. Had Austel tested the ballot against its other statutory obligations of consumer protection and promoting efficiency in the industry, it might have found a ballot in conflict with both.

Kevin Morgan is National Industrial Officer with the Communication Workers' Union.

Here's How Our Token Ring Products Stack Up To The Competition's.

Stack us up to any other major hub vendor, and see for yourself. Cabletron offers more Token Ring modules for the intelligent hub than the competition. Period.

But that's only the beginning.

Our Token Ring products are driven by Intel's high-performance i960 32-bit RISC processor. The competition's? 16 bits. Our active Token Ring modules provide up to 24 ports. The other guy's? Half that. And we have an integrated solution for bridging Ethernet, Token Ring and Wide Area Networks with a single module. Their solution? A stand-alone box.

Our products use Flash EEPROM for easy upgradeability. Theirs don't. Our Desktop Network Interface Cards support both STP and UTP and provide *at-a-glance* diagnostics. Theirs won't. And we've integrated twelve station ports on a single, manageable repeater module with Ring-In/Ring-Out capabilities for both fiber and copper. They can't.

With 19 modules for our MMAC hub, Cabletron's Token Ring solution is "the complete solution." Theirs? Well, they simply don't stack up.

For more information on Cabletron's Complete Token Ring Solution, give us a call at **(02) 878 5777.** And ask for your free copy of Cabletron's *SNA and Token Ring Integration Strategy,* the white paper that maps out a smooth migration to the "new" SNA and integrated SNA/Token Ring environments.

CABLETRON SYSTEMS Inc.

Token Ring

The Complete Networking Solution™

Unit 8, 112/118 Talavera Road • North Ryde NSW 2113 • (02) 878 5777

Facing the SNA Migration Dilemma

The demands on enterprise networks based on IBM's SNA are changing and network managers must now address a crucial migration issue — how to position the network for LAN interconnect.

SNA network managers are facing a dilemma: how to position their networks to accommodate the increasing volume of LAN-to-LAN traffic generated by the upsurge in peer-to-peer and client-server applications.

Most will have large investments in their SNA subarea network based on the typical hierarchy of VTAM hosts, local and remote communications controllers (37xx/NCP), and physical unit type 2.0 cluster controllers with their attached terminals and workstations. Even though much of the network traffic these days may come from Token Ring LAN-attached PCs running 3270 emulation software, that traffic is nevertheless hierarchical, with the PCs configured as dependent secondary logical units.

How then to handle LAN-to-LAN traffic, whether SNA or otherwise? Creating a separate LAN-interconnect network may be acceptable as an initial solution, if the LANs are local and the scale of LAN-to-LAN connectivity is small. But as the number of remote sites requiring interconnect grows, and the nature of the traffic becomes increasingly enterprise-critical, the issue comes to a head.

SNA network managers must adopt strategies which will enable their networks to support increasing levels of peer-to-peer and client-server connectivity, while at the same time avoiding undue disruption to existing applications. The fundamental choice is between evolving the hierarchical SNA network to IBM's Advanced Peer to Peer Networking (APPN), versus building an enterprise multiprotocol router network. The choice must primarily be driven by the business needs of the applications that will be networked and their resulting communications characteristics and protocol requirements. If these criteria allow you to adopt IBM's LU 6.2 peer-to-peer protocol as your primary application protocol, then a strong case can be made for evolving your network towards APPN.

The APPN Choice

APPN is back in the news. A few years ago one could be forgiven for thinking that it had lost the networking race, overshadowed as it was by the effusive reporting of LAN interconnect, the Internet, SNMP, TCP/IP and multiprotocol routers. But the recent effort IBM has put in to communicating its network plans (see 'IBM's Bid For Multiprotocol Nets,' *Australian Communications*, July 1992), the appearance of new APPN products like Networking Services/2, and the addition of APPN functionality to existing products like the 3174 and VTAM, are giving APPN a new lease of life.

APPN has some real advantages over the TCP/IP multiprotocol router alternative (see 'APPN Rises to the Enterprise Challenge,' *Australian Communications*, June 1993), not the least of which is the smooth evolutionary path that can be followed in getting from your hierarchical SNA network to your APPN network. The first step on that path is to LU 6.2. IBM's LU 6.2 (Logical Unit 6.2) offers the opportunity for the enterprise to move into peer-to-peer networking in an evolutionary way. Like NetBIOS and TCP/IP Sockets, LU 6.2 communicates between a pair of application programs rather than being tied down to physical devices.

A number of Australian companies have found it beneficial to initially gain application development experience with LU 6.2 while maintaining the traditional host-to-workstation hierarchy. With this approach, workstation-to-host applications are developed using LU 6.2, although the workstations are still configured as dependent LUs, relying on VTAM to mediate in session establishment. In concept they operate in a similar manner to the programmable logical units in IBM 4700 controllers, but rather than use an application defined LU 0 protocol, they use the LU 6.2 protocol.

If the networked applications are developed in-house and the conversational model is appropriate, then convincing your developers to adopt LU 6.2 should not be too difficult. It offers a rich set of peer-to-peer functions, including selection of transport characteristics (delay, cost, etc.), session and conversation level security, dynamic activation of partner programs, confirmation and synchronisation control, and extensive error notification. It is well documented, stable and proven. Furthermore, platform independence can be achieved by accessing LU 6.2 via IBM's Common Programming Interface for Communications (CPI-C), a single application programming interface available on OS/2, DOS and Windows PCs, AS 400s and S/390 platforms. For example, COBOL or C programs written to CPI-C could be developed on one of these platforms and ported to any of the others with little or no change.

Even if your networked applications are packages developed by third parties, you may be able to select LU 6.2-based products, as it is widely supported in the industry. In fact, LU 6.2 is fast becoming ubiquitous in SNA environments, so the risk of striking problems in the future due to lack of support skills will be minimised.

Peer-to-Peer Via the Backbone

Once the essentials of LU 6.2 have been mastered in a hierarchical environment, the next step in the evolution is to implement peer-to-peer communications at the session layer. This will allow session initiation by either partner, sessions with multiple partners concurrently, and parallel sessions. Note that we are not yet talking APPN. VTAM and NCP at the appropriate levels (VTAM 3.2, NCP 5.2 or NCP 4.3) support independent LUs and peer-to-peer communications at the session layer, all over a traditional SNA backbone (see Figure 1 on page 71). Existing LU 2 and LU 0 traffic will continue to be handled by the network as in the past.

Therefore, before making the step to APPN, you can gain experience with peer-to-peer sessions running over your SNA backbone. For example, if one of the partner applications in such a session is in an OS/2 PC running Communications Manager, the other partner could be an application on a host mainframe, or it could be an application in another OS/2 Communications Manager PC anywhere in the network.

In this latter case, each OS/2 Communications Manager views VTAM/NCP as a composite PU T2.1 low entry networking (LEN) node and has no awareness of the SNA backbone network. However, the BIND which is sent by the originating OS/2 Communications Manager to initiate a session with a partner LU will benefit from all VTAM's SNA directory services in locating the target OS/2 Communications Manager LU. Similarly, once the session is established, SNA routing, priority control and flow control will apply to the session data as it flows across the SNA backbone.

The final step is to make the transition to a full APPN network. From the mainframe perspective, VTAM Version 4 and the associated NCP will allow this transition to be staged according to your organisational requirements. Composite VTAM/NCP nodes will perform full subarea functions and full APPN functions, and provide the necessary transformations between the two networking modes at session setup time. This approach will allow parts of your subarea network to be converted to pure APPN piece by piece, as convenient.

At the LU 6.2 application and session levels little will change; peer-to-peer communications will be initiated and run in the same manner as in the LEN environment discussed above. The real impact and the benefits of APPN will be seen in terms of networking services. Benefits such as dynamic registration of LUs in End Nodes and Network Nodes, efficient directory services, automatic discovery of network topology and reaction to topology changes, optimal route selection, prioritisation based on class of service, and adaptive flow control are delivered by an APPN network.

What about 3270 sessions in an APPN network? IBM's direction is to include them as part of the APPN environment, so that 3270 sessions can also benefit from APPN networking services. IBM has not announced officially how this will be done. However, reports are that a 3270 LU's Network Node will act as a requester agent, and with the help of an intermediary VTAM Network Node server and the target application's Network Node, a session through the APPN network will be latched into place. Native LU 2 or LU 3 3270 data will flow over this session.

Although dependent LU 0 sessions could in theory also work in this way, in practice it would be prudent to migrate your LU 0 applications to LU 6.2. For LU 0 sessions to run natively in the APPN network, they would require a Network Node server with the requester agent software. The Network Node requester agent will apparently be initially implemented by IBM on the 3174 cluster controller. Whether it will be available more widely, and when, are moot points.

This sort of evolutionary approach to APPN will also position your enterprise network for the future networking developments foreshadowed by IBM: interconnection of multiple APPN networks, transparent bypass of failed links, lower level but faster routing options, rate based flow control, efficient utilisation of very high speed links, ATM and other fast packet-switching technology. These developments — APPN+ and Gigabit APPN — will push the architecture into gigabit and multi-gigabit packet technologies and bring enhancements in services and routing techniques.

LAN Protocols on an SNA Network

While moving towards APPN will cater for your SNA traffic, how will non-SNA LAN traffic be handled? This requirement will be driven by LAN workstation users wanting to access servers remote from their local LAN. If, for example, their LAN protocol is SPX/IPX, or NetBIOS, or TCP/IP — how can this be addressed with an SNA backbone and APPN architecture?

If the volume of such traffic is not excessive, then the immediate solution may be to use a product that transparently interconnects your LAN applications across your wide area SNA network. For example, products like Novell's SNA Links, IBM's LAN to LAN Wide Area Network program, Somerset Systems' Net Router, and Open Connect's IP Router.

Each of these products transports its native LAN protocol across an SNA backbone, using peer-to-peer APPC sessions across which LAN frames flow encapsulated in LU 6.2 request units. From the LAN user's perspective, no special action needs to be taken after the initial setup of the sessions across the SNA wide area network (WAN). The use of the application will be the same irrespective of whether the target server is on the local LAN or is remote.

Clearly, with LANs running at 4 to 16-Mbps, the performance impact of shipping data over WAN links of say 9.6Kbps will not be transparent to the user. So bulk file transfer is out. Nevertheless, for client applications accessing a remote server, where data volumes are relatively small, the products can offer acceptable WAN responsiveness. For example, central administration of remote file and print servers, remote database configuration, retrieval from e-mail post offices, and remote file access can all be ideal applications.

The challenge for all products of this type is to implement a design that efficiently transports the LAN data. Some products have been found wanting in this regard. For example, the Micro Tempus Enterprise Router, which utilised a host-based APPC

Figure 1: Steps on the Evolutionary Path to APPN

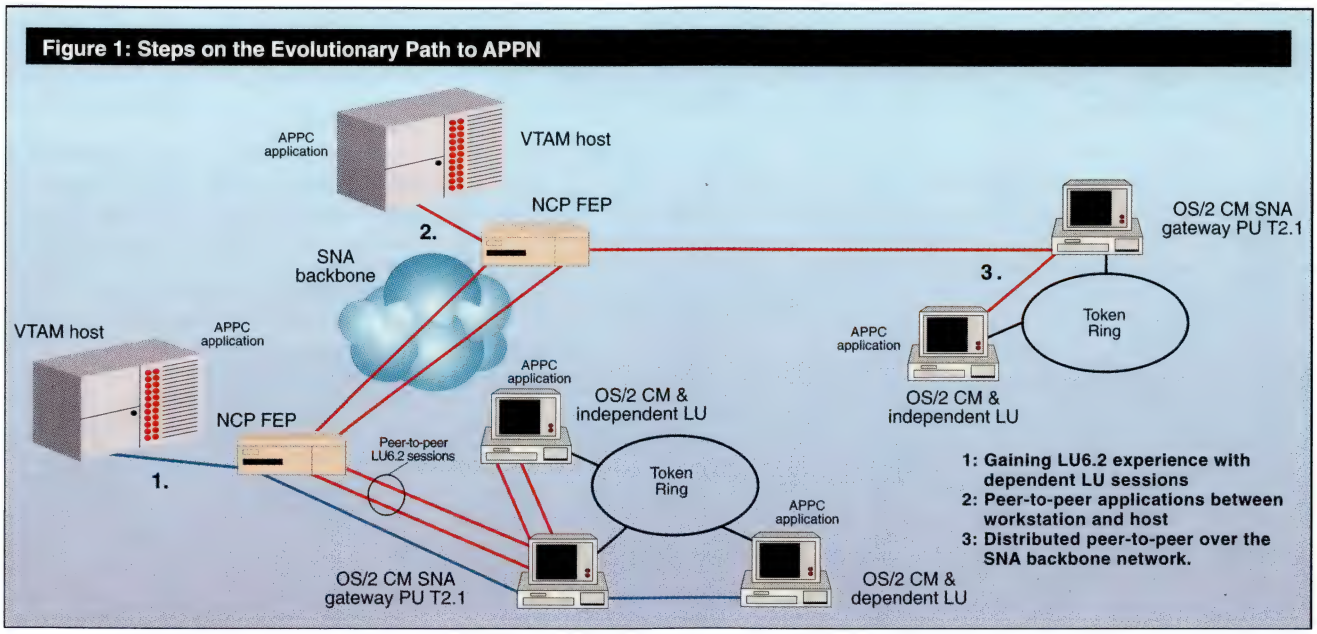

application to relay data flows from source LAN to destination LAN, just could not match the performance of the peer-to-peer products. Other design aspects each product must address include handling LAN broadcasts and timer constraints.

Unfortunately, none of the products offer any significant security features. Rather, they all rely on their LAN network operating system to provide security, an approach that is generally justified by the vendors as being in keeping with the objective of product transparency. For the potential applications indicated above, such an approach will often be acceptable. However, if remote databases contain live financial information or other enterprise critical data, business managers are likely to demand more in the way of access control and secure audit logging than the products currently provide.

SNA Links

SNA Links is the software product that allows NetWare LAN users to transparently communicate across an SNA network. It consists of two logical components: one for SNA encapsulation and one for IPX re-routing. Both components are present at the sending and receiving LANs, performing complementary roles. For example, at the sending end, the IPX router transparently re-routes NetWare data on the LAN to the SNA component, which encapsulates that data within the LU 6.2 protocol for transmission across the SNA WAN over a peer-to-peer session.

Users of SNA Links would be well advised to also implement Novell's set of WAN performance enhancement NLMs (NetWare Loadable Modules). SNA Links makes no provision to limit the forwarding of LAN broadcast traffic onto the WAN link, thereby exposing the link to the impact of the regular 60 second broadcasts issued by NetWare's Service Advertising Protocol. Novell's SAP Filter NLM addresses this problem, allowing you to limit the server advertising traffic flowing onto the WAN.

Secondly, to reduce the WAN impact of the NetWare Core Protocol's 'ping-pong' send-one-packet-acknowledge-one-packet traffic, the IPX Burst Mode NLM and corresponding workstation shell should be employed. Burst Mode allows a single file request to the server to result in the transfer of large data blocks (up to 64 Kbytes), between client and server. Through the use of a modified sliding window protocol, multiple packets can be transmitted in a burst sequence before an acknowledgment is required. A third Novell utility, the Large Internetwork Packet NLM, complements Burst Mode by supporting packet sizes larger than the usual IPX router limit of 576 bytes.

SNA Links runs as a NetWare Loadable Module under NetWare 3.11 and requires NetWare for SAA 1.2 or later to provide the LU 6.2 and PU T2.1 connectivity.

LAN to LAN WAN Program

The IBM LAN to LAN Wide Area Network program is an OS/2 product which interconnects local and remote LANs by routing NetBIOS sessions over an SNA WAN.

In concept, the LAN to LAN WAN program is similar to SNA Links, recognising and intercepting NetBIOS session activation and datagram frames addressed to one of its target stations. The program redirects the frames to the remote LAN, using peer-to-peer SNA LU 6.2 sessions established between itself and the remote LAN to LAN WAN program. NetBIOS applications, written to the IEEE 802.2 interface and running on user workstations on the local LAN, can then access remote servers using NetBIOS data flows. The product requires OS/2 Extended Edition 1.2 or later, with Communications Manager to provide the 802.2 and LU 6.2 support.

Although the IBM LAN Server network operating system is not a requirement, the other product requirements outlined above mean that in practice the LAN to LAN WAN product will only be found on 'true blue' LANs running LAN Server.

The LAN to LAN WAN program determines the source and destination stations it allows to communicate with each other, based on a list of NetBIOS names and/or station adaptor addresses configured with each copy of the program. This list is limited to 20 entries, but as wildcards can be used, the list can effectively be extended if your LAN nodes follow a structured naming convention. This name qualifier list is also the only mechanism for controlling the amount of NetBIOS broadcast traffic flowing onto the WAN — an approach which has the potential to lead to link utilisation problems if not actively managed.

To avoid NetBIOS timer problems across the WAN, the recommended approach with this product is to increase NetBIOS timer values within Communications Manager and to adjust LAN Server buffer sizes. Once again, careful tuning is called for to find the optimum balance in terms of LAN versus WAN performance. The product can send network management alerts to NetView or Net/Master, but rather than sending them directly, it relies on the services of IBM's LAN Network Manager.

The LAN to LAN WAN product was released in 1990, and while initially in the vanguard of these types of products, the current version is due for an update. IBM Australia advise that a new release is now available in the United States and will shortly be made available in Australia.

NetRouter

The other notable environment for NetBIOS session routing over an SNA network is between Microsoft LAN Manager LANs.

Figure 2a: The IBM 6611 Advantage

The Network Node capability of the IBM 6611 routers enables them to participate fully in the enterprise wide APPN network, in addition to providing multiprotocol routing functions.

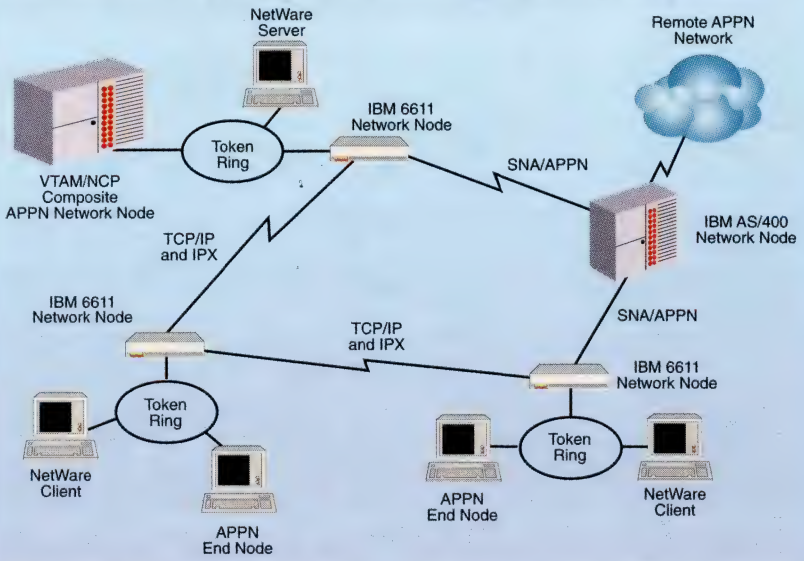

Figure 2b: The APPI Approach

The APPI Open Network Nodes will emulate APPN directory and routing services for their attached APPN End Nodes, in addition to providing multiprotocol routing functions. However, the APPI nodes cannot interwork with true APPN Network Nodes, so the APPI network will be logically isolated from the enterprise's APPN network.

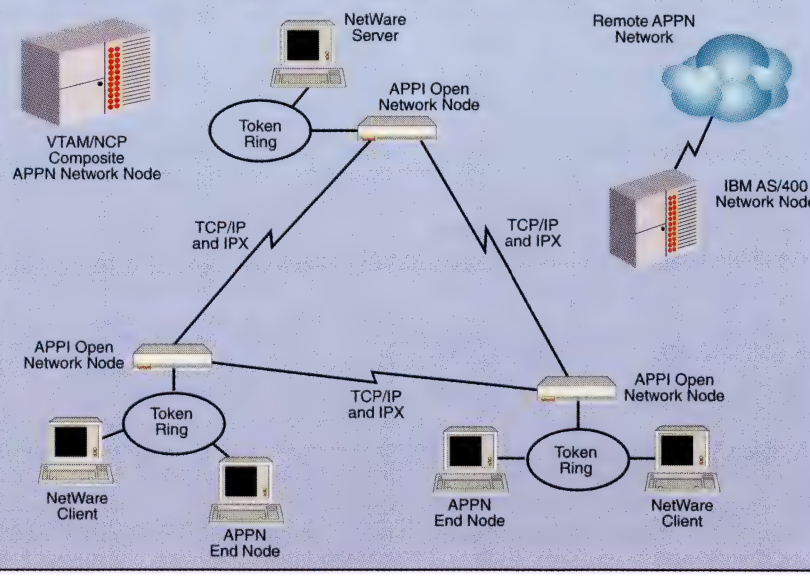

To address this need, Australian company Somerset Systems has produced NetRouter, an OS/2 product which interconnects LAN Manager or IBM LAN Server LANs.

Like SNA Links and the LAN to LAN WAN program, one copy of the NetRouter software resides on the source LAN, one on the destination. These copies would typically run on servers, although they can be on any OS/2 workstation. NetRouter dynamically loads the NetBIOS names of up to 225 remote servers or peer workstations and the associated remote NetRouter LU names.

When a user's workstation seeks to establish a NetBIOS session with a remote server the local NetRouter effectively acts as an agent for the remote device. The local NetRouter establishes parallel LU 6.2 sessions with the NetRouter at the remote destination. When the user workstation subsequently sends data on the NetBIOS session, the NetBIOS frames are transparently intercepted by the local NetRouter, encapsulated in LU 6.2 and forwarded over the SNA WAN to the remote NetRouter, for on-forwarding to the server on the remote LAN.

To minimise the impact of LAN-to-LAN traffic on WAN links, all NetBIOS broadcast traffic is filtered out by NetRouter. So that local LAN workstation users can identify remote servers accessible via NetRouter, each NetRouter does local spoofing of server broadcasts, generating server advertisements locally on behalf of servers on remote NetRouter LANs.

While some timer issues, such as session timeout, can be addressed simply by modifying LAN Manager parameters, the LAN Manager redirector presents more of a problem. Firstly, it has a fixed 30 second timer for completion of file reads and writes. Secondly, it defaults to a transmission option called Rawmode which sets up 64 Kbyte transfers for these reads and writes. While suitable for a LAN connection, these settings result in timeouts for such large transfers on WAN links of up to 19.2Kbps.

By setting a workstation initialisation parameter, Rawmode can be turned off and read/write transfers will be done in 4 Kbyte blocks. Unfortunately this applies to all transfers, whether on or off the local LAN, potentially impacting overall LAN performance. Somerset Systems have provided a convenient solution to this issue by providing a NetRouter command line utility, which the workstation user can use to dynamically turn Rawmode on or off as appropriate.

Like IBM's LAN to LAN WAN program, NetRouter can send alerts to NetView or Net/Master. However NetRouter does not require IBM's LAN Network Manager as an intermediary, but sends its alerts directly via the Common Service Verb interface of Communications Server or Communications Manager.

NetRouter requires LAN Manager 2.0, with Microsoft Select Communications Server 1.0 for LU 6.2 support. Alternatively, NetRouter requires IBM LAN Server 1.2 and OS/2 Extended Edition 1.2 or later, with Communications Manager providing the LU 6.2 support.

OpenConnect/IP Router

Distributed in Australia by Techway, the OpenConnect/IP Router is one of a range of products from OpenConnect (OC) Systems for interfacing TCP/IP and SNA networks. OpenConnect Server is the foundation product for this range and OC/IP Router builds on this foundation to encapsulate TCP/IP packets into LU 6.2 request units for transfer over an SNA backbone network to another OC/IP Router.

The OpenConnect Server can be implemented as a hardware or software product. The hardware version, on the SNA side, offers either Token Ring, SDLC or IBM 370 channel attachment. On the TCP/IP side, attachment is via an Ethernet or 802.3 connection. Once configured and operational, this hardware version can be considered as a 'black box,' requiring no user support. The software version of the OpenConnect Server is Unix-based and runs on Sun, RS/6000 and other common Unix platforms.

Both versions are functionally similar, providing the capabilities of a protocol con-

Why play games with your communications system

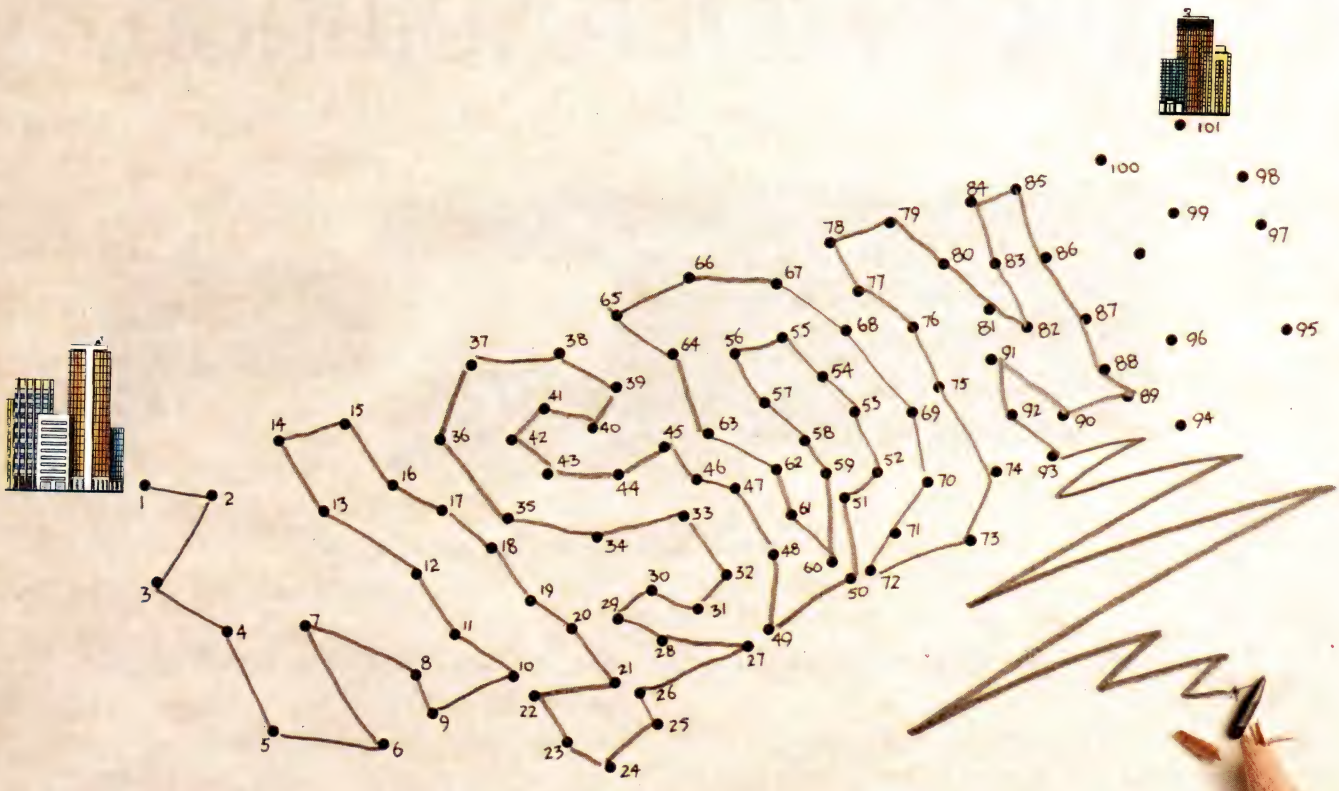

when there could be a better path to follow...

... Q-Net Australia.

For most business users, the current carrier services can present a torturous and sometimes costly path to completion. As a facilities manager we understand the strengths and weaknesses of the Carriers products, the available hardware options and service delivery.

While we compete against the main carriers and hardware vendors, we also use the best that they and others have to offer.

Through our expertise, we are able to offer our clients the best combination of technologies and deliver 15% to 25% savings on their current network costs – with no capital outlay.

If this sounds too good to be true, allow us to conduct a no obligation audit of your current system. We'll prove the benefit of selecting the most appropriate paths to follow.

Q-NET

AUSTRALIA

POINTING YOU IN THE RIGHT DIRECTION

BRISBANE (07) 852 1414 SYDNEY (02) 267 4884 MELBOURNE (03) 526 3506

GCA4817

Think of this as the match that lit the fuse that started the FDDI revolution.

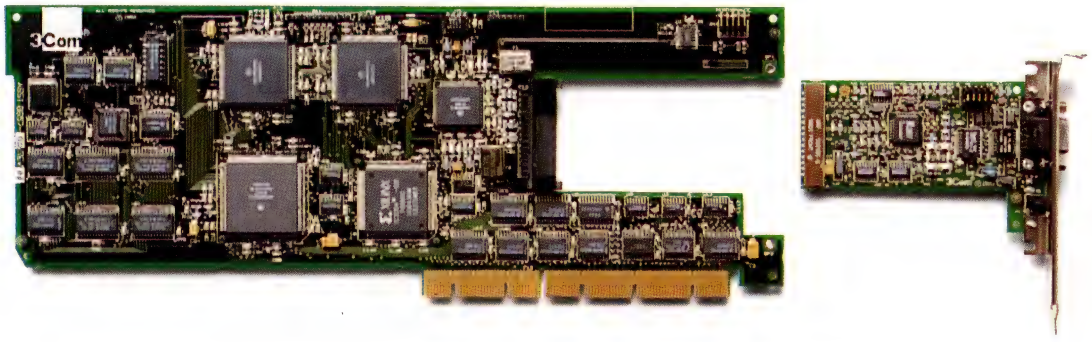

FDDILink's unique "PMD" module design makes changing media a snap.

Every revolution has its moment of truth. For FDDI, that moment is now.

Introducing 3Com's FDDILink™ family of 32-bit EISA adapters. They bring blistering FDDI performance to your network for as little as $2,500. That's up to half the cost of comparable adapters.

And that's not all. Backed by a lifetime warranty, the FDDILink family gives you the flexibility to ease your network into FDDI, at your own pace. Thanks to their unique "PMD" module design, FDDILink adapters can be used with shielded twisted-pair cable or fiber-optic cable. And even with unshielded twisted-pair when standards are finally approved. Simply plug in the appropriate media module.

Now you can let the high-powered hardware and software on your network perform at its full potential. That's because FDDILink can deliver 100 Mbps performance to the desktop using lower-cost shielded twisted-pair copper cable and connectors.

And what's even better, when you add 3Com's NETBuilder® II Routers and LinkBuilder® 3GH Hubs, you have a total solution that runs all the way from the desktop to the WAN.

But that's enough talk. For FDDI revolutionaries, this is a time for action. Call **(02) 959 3020, ext. 22**, or contact your nearest 3Com distributor.

3Com®

Networks That Go the Distance™

Tech Pacific Australia
(02) 697 8666

Tech Pacific New Zealand
(09) 415 8300

MM Data Networks Australia
(02) 980 6922

BICC Communications New Zealand
(09) 358 0322

i.t. connXions Australia
(02) 415 0555

verter, a gateway and a router in one system. If you planned to utilise a number of the OpenConnect applications, such as TN3270 or TN5250 terminal emulation, FTP file transfer, and e-mail bridging, or if you planned for a large number of concurrent sessions, the hardware OpenConnect Server may be appropriate. For lower levels of demand, running the software OpenConnect Server in the background on an existing Unix host may be suitable.

The OpenConnect/IP Router consists of routing software that resides with and utilises the OpenConnect Server. OC/IP Router versions are available for both SNA PU 2.0 and PU 2.1 connections. The former relies on a control program on the IBM MVS host to do re-direction from source to destination. The latter utilises peer-to-peer LU 6.2 sessions and avoids host involvement in data transfer. As mentioned previously, this latter approach is preferred from a performance perspective.

To assist in OC/IP Router configuration, OpenConnect offers the capability for centralised route management via either NetView or Net/Master. Configuration essentially requires the mapping of each source IP address to an SNA local LU/remote LU/mode combination.

Broadcasts in an Internet Protocol network are predominantly limited broadcasts, remaining on the local physical LAN. This, and the selective forwarding done by the router component of the OC/IP Router, means that the problem of LAN broadcasts flooding the SNA WAN links is avoided. Similarly, timeout problems caused by delays on the WAN links are handled by TCP's adaptive retransmission algorithms that adjust to varying delays and congestion.

LAN Protocols on an APPN Network

The use of a software package to route LAN interconnect traffic over your SNA/APPN backbone may be a good interim solution. It may even be regarded as a longer term solution for those LANs that service just a handful of users, where the LANs are currently multi-dropped off a low speed line and the costs of upgrading to point-to-point links for bridge/router connectivity cannot be justified.

Nevertheless, in your regional offices and other major corporate centres, it is likely that the demands that your users will place on the network for LAN interconnect will grow to outstrip the capabilities of the SNA/APPN encapsulation approach. How then do you handle LAN protocol traffic from this group of users with your APPN network?

One approach will be to configure corporate centres into a number of subnetworks based on multiprotocol routers with APPN Network Node capability. At the LAN interface, in addition to supporting protocols such as SPX/IPX, NetBIOS and TCP/IP, such a router will also present an APPN Network Node server interface to APPN End Nodes. So the End Nodes on the LAN will be able to dynamically register their LUs with the router, submit directory search requests for APPN LUs, receive APPN session route information and LU search requests from the router, and establish SNA/APPN sessions.

For companies that do not want to be dependent on a single supplier, this approach currently carries some risk, as the IBM 6611 is the only router which currently has APPN Network Node capability. However other router vendors such as 3Com and

APPN Connection Networks

IBM has coined the term 'Connection Network' to refer to a group of Token Ring attached APPN End Nodes and Network Nodes which can exchange session data over direct logical links. It is true that establishing a logical link connection between a pair of stations on a LAN is nothing unusual in most LAN environments, however, if we look more closely at APPN we can see why connection networks are necessary.

APPN routing tables are dynamically established to reflect the possible routes between all Network Nodes. When one End Node requests session establishment with a remote End Node, the resulting session path is over the local End Node-to-local Network Node connection, then over the optimal APPN backbone route, and finally over the Network Node-to-End Node connection at the remote end. Typically then, the session path will be via one or more Network Nodes.

What happens if the two End Nodes are attached to a common Token Ring LAN? Unless a logical link is defined between them, the resulting session path would still be via their serving Network Node, as this would be the only path known to the APPN network. To avoid this situation, at each End Node we would need to define logical links to every other End Node on the LAN. Such a requirement would clearly be a backward step in the quest to simplify APPN configuration. Consequently, the connection network concept was developed.

A group of APPN nodes can be defined as belonging to the same connection network if each node's LAN adaptor can access any of the other adaptors in the group. When an End Node requests session establishment with another member of the connection network, the APPN Network Node will advise the optimal route, and will include the target End Node's LAN adaptor address. The initiating End Node can then dynamically establish a LAN logical link to the target, allowing session data to flow directly rather than via the Network Node.

Peter Johnson

APPN Connection Network Definition

Without an APPN connection network definition, data on a session between two End Nodes on the same local area network will flow through the serving Network Node, creating additional LAN and Network Node congestion. With a connection network, data flows directly between the End Nodes.

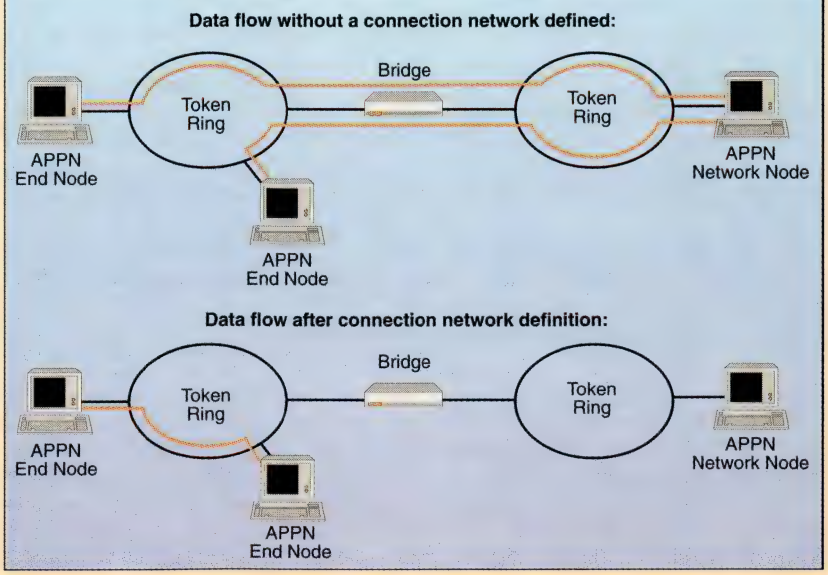

The APPI Challenge

Advanced Peer-to-Peer Internetworking (APPI) is the result of a Cisco Systems initiative which appears to be gaining significant support. Announced at Interop in San Francisco last October, the Cisco-led APPI Forum now includes such companies as Digital Equipment Corporation, Hewlett-Packard, Unisys, SynOptics, Cabletron Systems and Proteon. The Forum's objectives include:

- Development of an open, TCP/IP-oriented solution for SNA peer-to-peer networking;
- Production of an APPI specification that will be available, free of charge, to vendors, users and other interested parties; and
- Promotion of industry cooperation to achieve growth in the use of APPI products and services. An example of such cooperation is the proposed establishment of an interoperability testing laboratory.

To APPN End Nodes, an APPI router (designated an Open Network Node, ONN) will appear to provide standard APPN directory, routing and session establishment services. However, unknown to the End Nodes, the backbone or wide area network that the ONN is connected to will be a TCP/IP network.

In addition to the attraction of open standards, the APPI Forum claims two major benefits. On the technical side, the APPI approach offers adaptive re-routing to recover from path failure, more extensive media support, and avoids adding another protocol to the backbone. On the commercial side, network managers will not need to replace their existing TCP/IP networks to achieve SNA peer-to-peer networking.

The APPI specification is expected to be submitted to the Internet Engineering Task Force soon, with official approval likely to take about two years. The first product demonstrations are targeted for the third quarter this year, although Cisco demonstrated a working prototype at the Interop Spring show in Washington in April. In the demonstration, APPI's Open Network Node (ONN) software ran on an IBM RS/6000 workstation, with Cisco routers providing the TCP/IP transport network and the APPI distributed directory services. PCs running OS/2 Extended Services acted as SNA end systems, while RS/6000s and Cisco routers delivered the APPI networking functions. Cisco says the prototype fully implemented SNA LU 6.2 and Node Type 2.1, with the ONNs supporting LU 6.2 communication between the SNA peer end systems represented by the PCs.

So when considering APPN, how are APPI routers likely to stack up against IBM's 6611 router?

The major deficiency of APPI routers will be that they cannot interwork with true APPN network nodes. The IBM 6611 on the other hand, can participate fully in an APPN backbone network, providing intermediate session routing to adjacent Network Nodes as well as all the other APPN functions — directory services, topology and routing services, and so on.

The major strength of APPI routers is likely to be in performance. When considering SNA/APPN traffic between a pair of routers, both APPI routers and IBM 6611 routers encapsulate that traffic in TCP/IP packets for transport across the wide area network. Both APPI routers and IBM 6611 routers have their multiprotocol routing functions to do, including maintaining topology awareness. However the IBM 6611 effectively has the burden of duplicating many of these functions, providing them both for the multiprotocol environment and for the APPN environment.

Peter Johnson

Novell have licensed APPN Network Node code from IBM and intend to deliver APPN products. As APPN garners an increasing user base, we can expect other router vendors to join in. Both Cisco and Proteon, for example, have indicated they will offer APPN Network Node capability in 1994.

Note that the Advanced Peer-to-Peer Internetworking (APPI) concept introduced by Cisco Systems in August 1992 only partially satisfies the above APPN router scenario (see 'The APPI Challenge'). Although APPI routers will appear as Network Nodes to the End Nodes on their LAN interfaces, they do not provide APPN Network Node support at the WAN interface. Rather, they use TCP/IP and existing router protocols like RIP and OSPF. If your peer-to-peer SNA applications are all on the periphery of your enterprise network, located in APPN End Nodes, then APPI routers would satisfy your needs. If, however, you have other true APPN Network Nodes (and many enterprise nets will have at least one VTAM/NCP Composite Network Node), then routers with Network Node capability will be required for full internetworking (see Figure 2 on page 66).

The other approach to handling the increasing levels of LAN protocol traffic from major corporate centres, along with your APPN network, is to implement separate logical LAN interconnect and APPN networks on a common transmission network.

However, this approach has potential disadvantages when compared to the integrated APPN router network approach discussed above.

Firstly, equipment costs may be higher. Maintaining both APPN Network Nodes and multiprotocol routers not only duplicates network routing equipment but also requires equipment to multiplex the two logical data streams. The common transmission network between corporate centres is likely to require links in the range of 128Kbps to 2Mbps, depending on the nature of the applications and their traffic volumes. You will be faced with the cost of multiplexers, bandwidth managers or frame relay switches to overlay the logical networks onto transmission services such as ISDN or Megalink. Similarly, if you use Flexnet, the cost of the Unimux multiplexers is effectively built into the tariffs.

Another service choice for implementing a common LAN interconnect transmission network is Fastpac 2, which became available in April this year. LANs containing APPN nodes could utilise the Fastpac 802.5 Token Ring interface option and retain APPN features by using the APPN connection network facility. This facility enables an End Node to discover the MAC address of its desired partner End Node and dynamically establish a logical link to it across the LAN (see 'APPN Connection Networks' on page 69). Routers can also attach to Fastpac 2 using the Token Ring interface option. Reports are that the use of transparent bridging rather than source route bridging (SRB) on this interface has on occasions created some performance problems, however these will be overcome with the release of the Fastpac 2 SRB option later this year.

The second potential disadvantage of separate logical LAN interconnect and APPN networks compared to the integrated APPN router network approach is greater network management complexity. Two logical networks increases configuration demands, adds another dimension to fault diagnosis, and makes it even more difficult to obtain support staff with the full suite of skills.

Despite these potential drawbacks, some network managers still find that the approach of maintaining separate SNA/APPN and multiprotocol router networks has its place. Where a corporate voice-data transmission network is already established, or is planned, the incremental cost of allocating additional LAN interconnect channels will be small. For some, this approach also has the additional benefit of keeping open LAN interconnect options in uncertain times.

Peter Johnson is a Senior Consultant at Communications Network Management Pty Ltd (Sydney) where he specialises in data communications and network management.

Why the London Stock Exchange chose Ericsson Business Network.

Being able to communicate is of vital importance. Especially when you turn over hundreds of millions of pounds every day. That's why the International Stock Exchange in London installed an Ericsson Business Network.

The Exchange needed a communications system that offered the best in efficiency, reliability and security. Business Network answered all the Exchange's voice and data needs, providing the perfect solution for both now and in the future.

The Exchange's data network uses an Ericsson ERIPAX packet-switching system. Its administrative telephone system is co-ordinated through an Ericsson MD110 digital PABX.

Right now Ericsson are delivering similar solutions for a large number of Australian companies. If you'd like a free Business Network information kit giving details, call Ericsson.

*business*network

THE COMPLETE COMMUNICATIONS NETWORK

Vic. (03) 480 4888 N.S.W (02) 805 4999 QLD (07) 262 6222 S.A (08) 234 5555 W.A (09) 262 6222 A.C.T (06) 257 1033 Tas (003) 31 2544 P.N.G (675) 25 6566

ERICSSON

Sanford Vick 3428

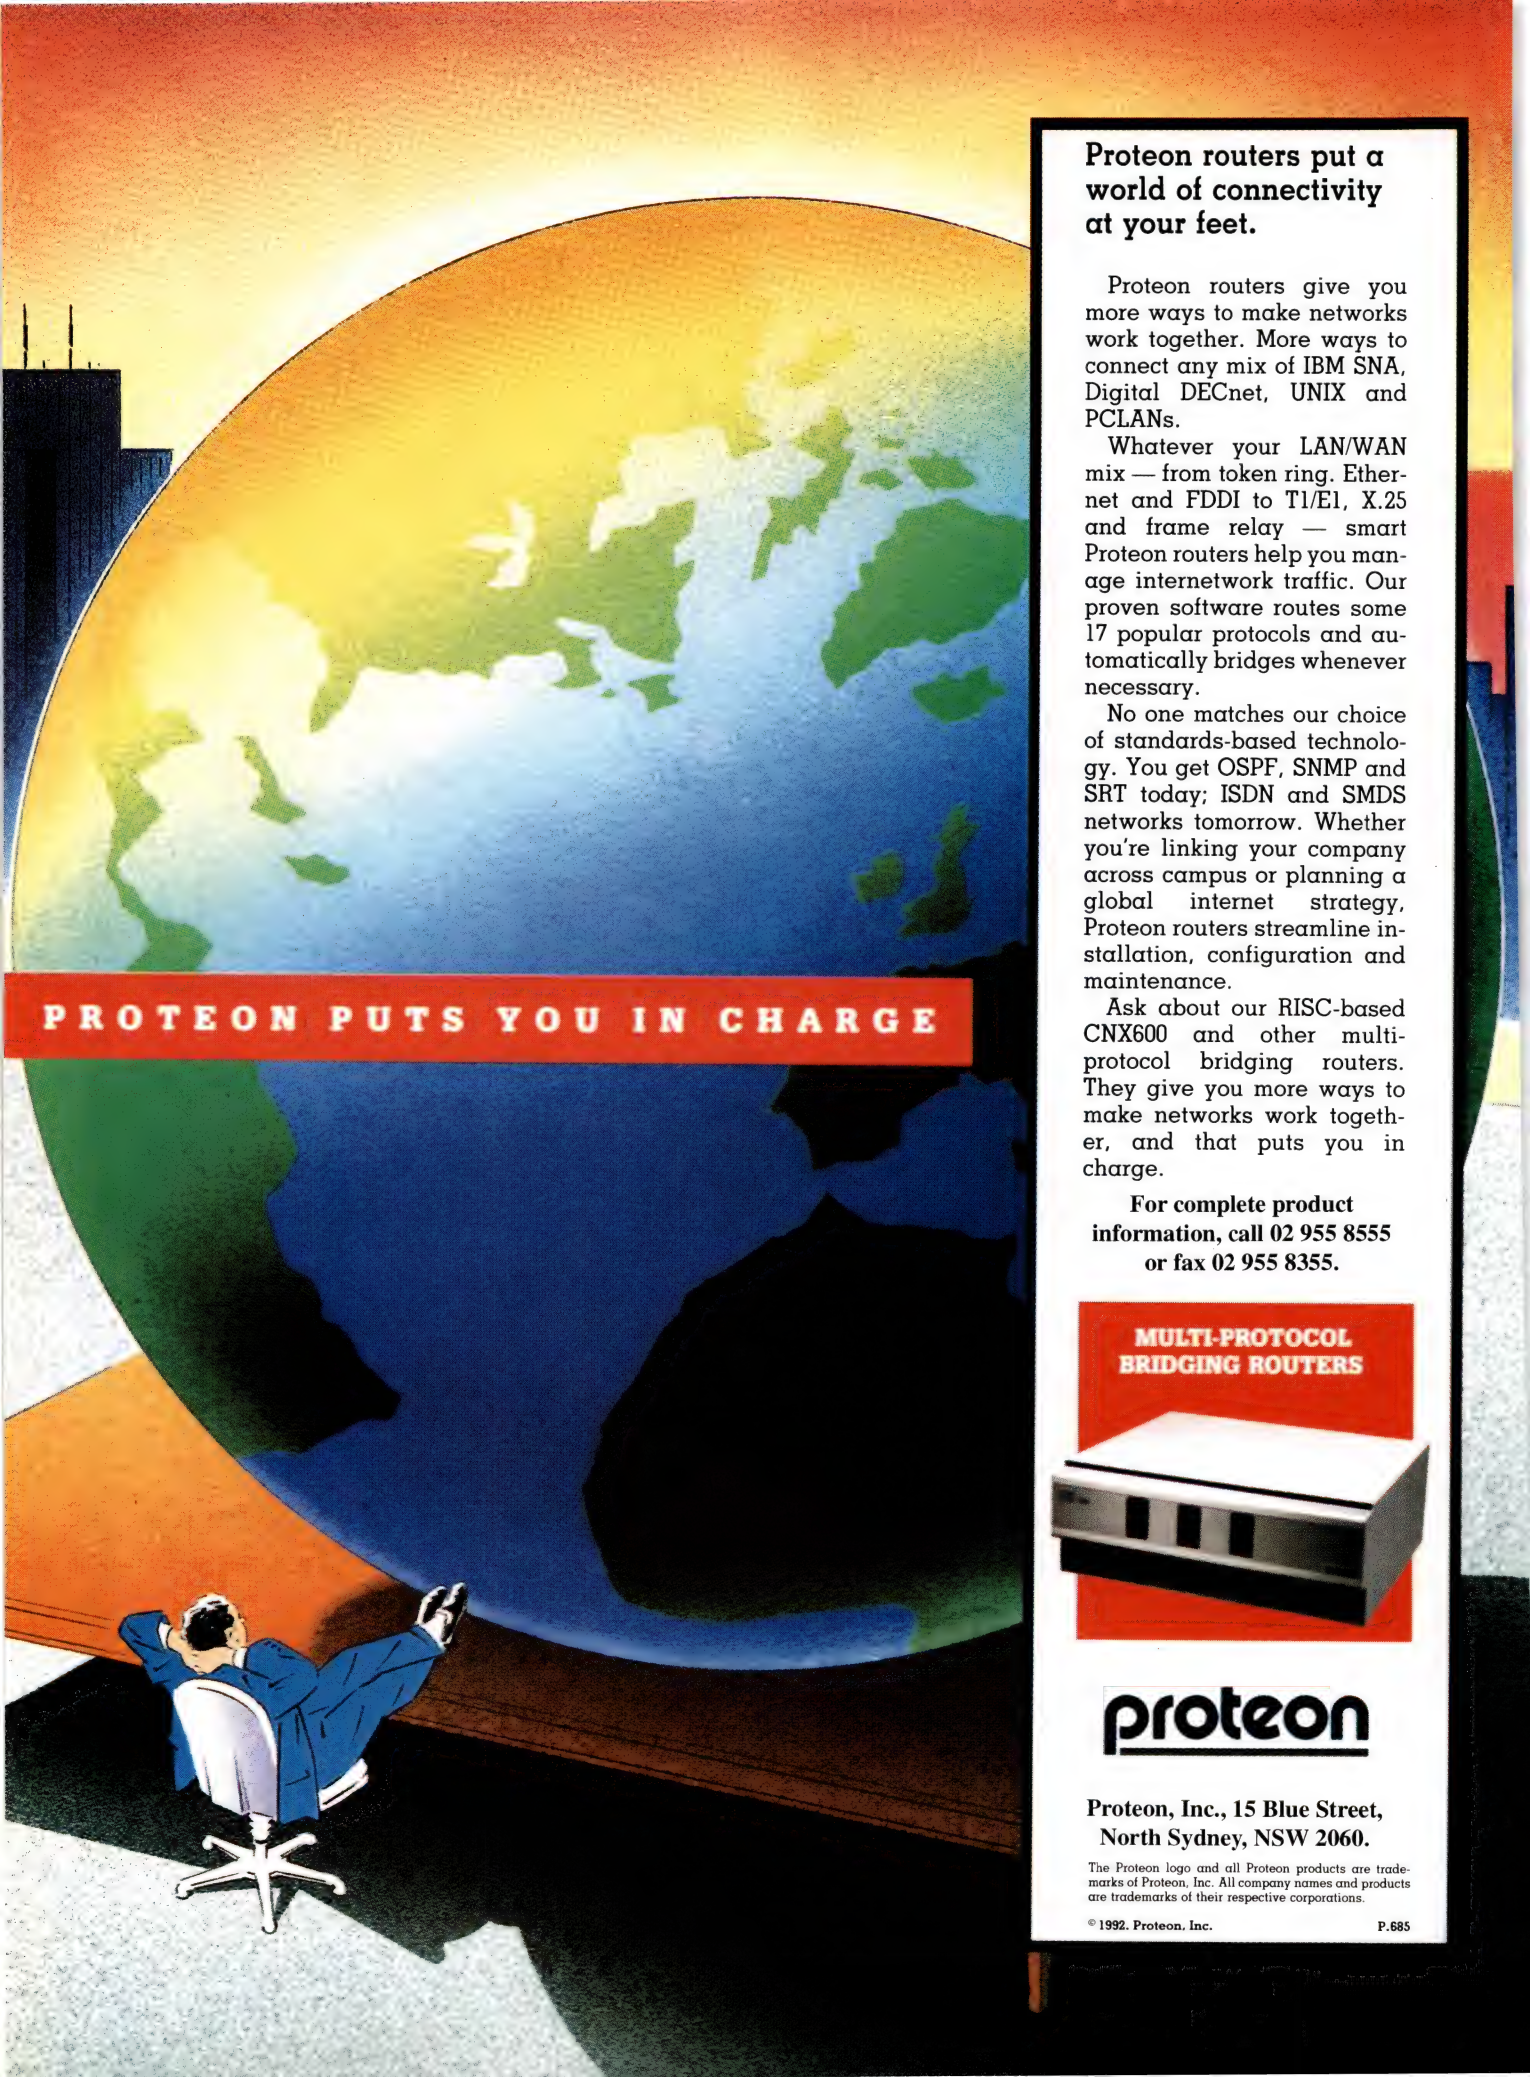

Proteon routers put a world of connectivity at your feet.

Proteon routers give you more ways to make networks work together. More ways to connect any mix of IBM SNA, Digital DECnet, UNIX and PCLANs.

Whatever your LAN/WAN mix — from token ring. Ethernet and FDDI to T1/E1, X.25 and frame relay — smart Proteon routers help you manage internetwork traffic. Our proven software routes some 17 popular protocols and automatically bridges whenever necessary.

No one matches our choice of standards-based technology. You get OSPF, SNMP and SRT today; ISDN and SMDS networks tomorrow. Whether you're linking your company across campus or planning a global internet strategy, Proteon routers streamline installation, configuration and maintenance.

Ask about our RISC-based CNX600 and other multiprotocol bridging routers. They give you more ways to make networks work together, and that puts you in charge.

For complete product information, call 02 955 8555 or fax 02 955 8355.

proteon

Proteon, Inc., 15 Blue Street, North Sydney, NSW 2060.

Kevin Tolly and
Eric Hindin

Can Routers Be Trusted With Critical Data?

Can lethally congested router internetworks be trusted with mission-critical applications? For most network managers, that's a question of life and death.

When it comes to mission-critical applications, getting high-priority data across a corporate network isn't the main thing, it's the only thing. No big deal, right? Unless, of course, the corporate network in question relies on bridge/routers to move data from Point A to Point B. In that case, even the simplest business transaction can turn into a hit-or-miss proposition.

It's no secret that bridge/router internetworks are starting to sag — and sometimes collapse — under the crushing burden they're asked to bear. Complex, competing applications; so-called unroutable protocols (like SNA and NetBIOS); and traffic pileups on low-bandwidth WAN links are pushing congestion to lethal limits — and beyond. It should come as no surprise, then, that network managers are leery about risking critical applications on already overburdened internets, playing Russian roulette with vital data (and their careers).

Router vendors claim to have the answer. In the past six months or so, they've delivered a bewildering mix of software enhancements intended to address reliability problems. Protocol prioritisation, broadcast reduction, name-caching, LLC2 local acknowledgment, enhanced dynamic rerouting, tunnelling, discovery assist — the list certainly sounds impressive. And when the vendors tell it, there's no doubt that their particular offering is just the thing to solve all the problems on that troubled internetwork. Maybe.

In the industry's first-ever evaluation of these recent releases, the US-based Data Comm Test Lab took a close look at the one feature that virtually all vendors say will boost network reliability: protocol prioritisation. (Other features will be checked out in later tests.) Protocol prioritisation is a simple concept: network managers assign a priority to one or more protocols so that they'll be guaranteed more bandwidth than other protocols transmitted over the same physical link. Hence, high-priority data will always get where it's going.

The Lab tested routers from Advanced Computer Communications, Cisco Systems, Crosscomm, RAD Network Devices, and Wellfleet Communications. Along the way, it made some striking discoveries and came to one overall conclusion. Despite vendor assurances — and the very real promise of prioritisation — bridge/router internetworks still can't deliver the stability of SNA installations. Prioritisation schemes range widely in sophistication and effectiveness. Most were difficult to configure. And once prioritisation was configured, judging its effectiveness proved extremely difficult. In some cases, turning prioritisation on made throughput take a noticeable dip — exactly the opposite of what the Lab expected to see. In other instances, enabling prioritisation threw stable parameters out of whack, which called for tedious, time-consuming reconfigurations.

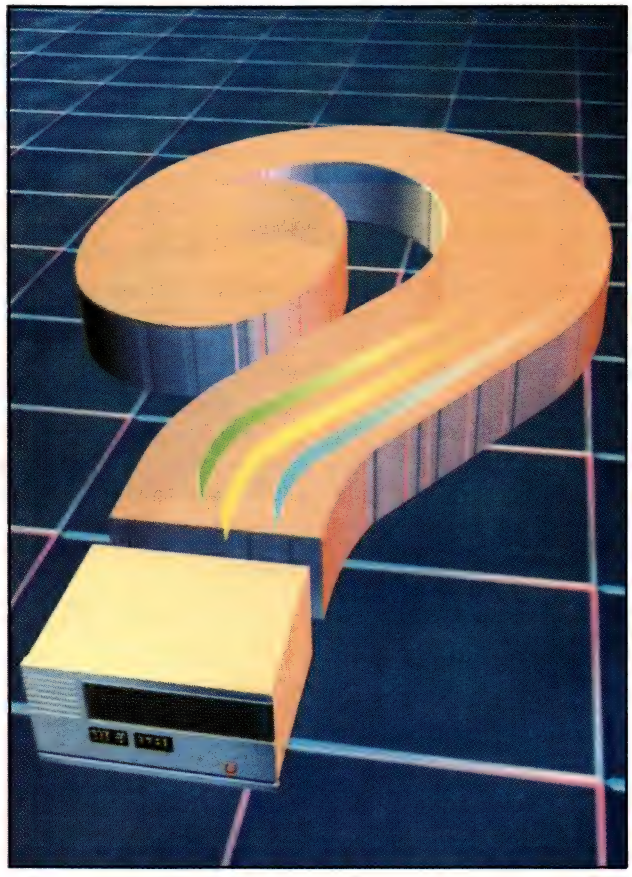

The fault may not lie entirely with the vendors. Protocol prioritisation is new. So new, in fact, that it was impossible for the Lab to determine if it had fully optimised each vendor's product. And in most cases, the technical support personnel that visited the Lab to help with the tests were not yet trained to use prioritisation. Thus, the Lab's findings should be treated as indicative of what a par-

Routers on Review

VENDOR	PRODUCT	SOFTWARE RELEASE	ROUTED PROTOCOLS	BRIDGING PROTOCOLS
Advanced Computer Communications Ungermann-Bass (03) 696 2006	ACS 4200	6	TCP/IP, IPX, XNS, DECnet, AppleTalk	Source routing, spanning tree
Cisco Systems (02) 957 4944	Cisco 4000	9.1.1	TCP/IP, IPX, XNS, DECnet, AppleTalk, Vines, SDLC transport	Source routing, spanning tree
Crosscom Scitec Communication Systems (02) 428 9555	ILAN Universal Router	5.04.20	Protocol Independent Routing (PIR), TCP/IP, IPX	PIR supports source routing and spanning tree traffic
IBM (02) 634 9111	6611 Models 140 and 170	1.1.0.1	TCP/IP, IPX, XNS, DECnet, AppleTalk	Source routing
Proteon (02) 955 8555	CNX 500	13	TCP/IP, IPX, XNS, DECnet, AppleTalk, Vines	Source routing, spanning tree
RAD Network Devices Dataplex (03) 210 3333	RTB-40	4.1	Routes protocols using proprietary approach based on SPF	Proprietary approach supports source routing and spanning tree
Wellfleet Communications (02) 959 1290	Backbone Link Node	7.5	TCP/IP, IPX, XNS, DECnet, AppleTalk, Vines	Source routing, spanning tree

IPX = Internet Packet Exchange SPF = Shortest Path First SDLC = Synchronous Data Link Control XNS = Xerox Network Services

ticular bridge/router can do, rather than as a hard-and-fast gauge of performance.

It also needs to be noted that all the priority-setting schemes now being brought to market are absolutely proprietary. Network managers are thus forewarned: as soon as prioritisation is enabled, multivendor interoperability is lost.

Given all these difficulties, the throughput figures presented should not be taken as exact measures of performance. Simulating the conditions that cause prioritisations to kick in and predicting behaviour as network loading changes are at this point more black art than hard science.

What the Lab can do is give network managers an advance look at some of the difficulties they will face when trying to choose the latest bridge/routers. Prioritisation is an attempt to deal with a very real problem, but until vendors can back up their claims with stable, reliable products, it's wise to proceed with caution.

Congested Arteries

All prioritisation schemes — regardless of underlying algorithms or fancy names — have the same principal purpose: to maintain session integrity. Simply put, they give network managers a way to ensure access to specific applications with guaranteed response times — no matter how heavily loaded the network becomes.

That's a critical guarantee, since congestion on many corporate networks is reaching the point at which sessions can't get through. This sort of lethal overloading can have myriad causes: it may be that the number of communicating nodes is so high that overhead alone consumes an inordinate amount of bandwidth or that unpredictable traffic patterns make it near impossible to allocate resources. It may be that 2Mbps pipes are choked with protocols intended for 10Mbps LANs or that bargain-rate 48Kbps WAN links to remote sites are backed up with huge file transfers.

Most likely, some combination of all of the above is threatening the stability of the internetwork, which makes running critical applications a gamble that most net managers won't want to chance. Prioritisation is one way to help safeguard vital data.

In essence, prioritisation can be defined as any scheme that alters the first-in, first-out approach that bridge/routers typically take with data (see Figure 1). Prioritisation, in contrast, is actually a multi-stage process. Incoming packets are first examined in order to determine which protocol is present. They're then queued for transmission by the prioritisation algorithm.

All of the vendors tested use their products' filtering capabilities to determine which protocols are used. (Filters are typically used to examine packets and prevent unnecessary traffic from travelling across the link.) The whole notion of filtering tends to be fairly primitive, since it depends upon searching every bit in a packet in real time until the target field is found. Specifically, the prioritisation filters now in use

Figure 1: First In, First Out

Without prioritisation, routers handle traffic on a straight first-in, first-out basis, which can easily result in high-priority data being blocked and mission-critical applications being dropped in favour of less important sessions.

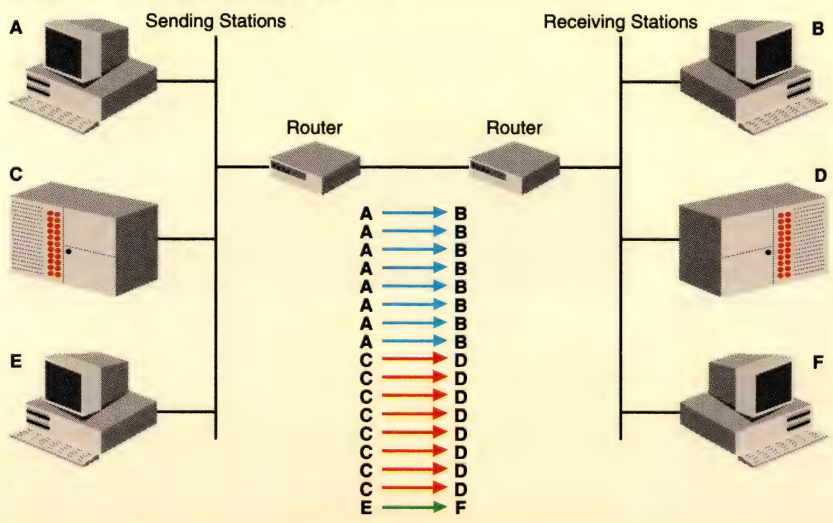

MAXIMUM WAN PORTS	PRIORITISATION SCHEME
2	Bandwidth sharing, with priorities
4	Priority queuing
5	Priority queuing
12	Data Link Switching
4	Bankwidth sharing, with priorities within protocols
4	Priority queuing
16	Priority queuing

determine which protocol a given frame carries by looking at the 8-bit Service Access Point (SAP) field, which contains the Layer 2 address and several bits that indicate the protocol being used.

The trouble here is that the SAP is not a foolproof indicator of which protocol the packet carries. The values the SAP now uses to indicate protocols are based upon convention. F0, for instance, is typically used for NetBIOS traffic, but nothing prevents an application vendor or end-user from assigning A0 to NetBIOS (as long as both sides of the session use the same assignment). That would make it impossible for a router to identify A0 packets as NetBIOS traffic.

Setting Priorities

There's a larger issue: filtering based on SAPs can't tell the difference between applications that use the same protocol. In other words, the bridge/router can't tell the difference between, say, NetBIOS traffic sent from Lotus Notes users and NetBIOS traffic sent from an IBM LAN server — even though users see them as independent applications with very different priorities.

One or more algorithms are used to set the priorities themselves. Although vendors have come up with a range of proprietary schemes, there are really only two underlying algorithms. They may be used alone or in tandem.

The first is based on the concept of precedence. In effect, it gives one protocol precedence over another, thus allowing it to be transmitted first. This approach is sometimes referred to as precedence queuing, because it's generally implemented with queues (areas of memory set aside for a specific purpose). When a packet carrying a high-priority protocol enters a router, it's sent to the high-priority queue. Other priorities also can be established, along with their corresponding queues. The router transmits all data from the high-priority queue before sending any other packets onto the network (see Figure 2). Cisco, Crosscomm, RND and Wellfleet are taking this approach.

There's a problem, though. If enough high-priority traffic enters the router, the high-priority queue will never be emptied. Thus, none of the other queues will ever be serviced and all lower-priority traffic will be blocked.

There's no surefire solution. The vendors that use precedence queuing allow the size of the high-priority queue to be adjusted. They also make it possible to adjust the amount of time that packets are held within queues (latency). Both of these schemes should help if high-priority traffic is bursty: Part of the burst will be transmitted and the rest discarded when the high-priority queue overflows. In the time it takes end-stations to request retransmission, the router should be able to service the other queues. But this is hardly an elegant solution, and again it's of little use if the high-priority queue is never emptied.

The second queuing algorithm is based upon the concept of shared bandwidth. A portion of the available bandwidth is reserved for every protocol, with high-priority protocols getting the largest share. A shared-bandwidth scheme thus never allows any protocols to be completely blocked.

There are two types of sharing algorithms: static and dynamic. In the first, the net manager manually sets the amount of bandwidth available to each protocol. SNA, for instance, might get 65%, with the remainder divided up among other protocols. Proteon takes this approach. What's more, the vendor's software automatically allows bandwidth to be borrowed from underutilised segments when necessary.

Dynamic-sharing algorithms use the router to continuously monitor traffic flow and allocate bandwidth according to need. ACC is taking this approach. Its ACS 4200 keeps tabs on the number of ongoing sessions and makes sure that every session is regularly serviced (see Figure 3 on page 80). Since a dynamic sharing algorithm handles each session individually, there's no chance that a conversation will be terminated accidentally because another application is hogging more than its fair share of bandwidth (exactly the sort of problem that occurs with conventional bridge/routers on congested networks). Further, this sort of algorithm can be augmented by specifying explicit priorities for protocols. For example, if one protocol is designated high priority, two packets could be transmitted for every packet of a lower-priority protocol.

The underlying algorithms used to set priorities are only a very rough gauge of a product's overall effectiveness. Other issues must be taken into account. For instance, although vendors are quick to claim their products can set priorities for all supported protocols, this statement can be somewhat misleading. Every vendor's product literature indicates that its algorithm can prioritise SNA traffic. What many mean is that they can prioritise source route bridging traffic.

The trouble is that protocols besides SNA use source route bridging. Cisco and Wellfleet group SNA in with other bridged pro-

Continued on page 80

Figure 2: Giving Traffic the VIP Treatment

One way to ensure that mission-critical data doesn't get stalled behind packets carrying less vital information is with a precedence-queuing algorithm. This enables net managers to assign priorities to protocols and makes sure the router services the highest-priority sessions ahead of any others.

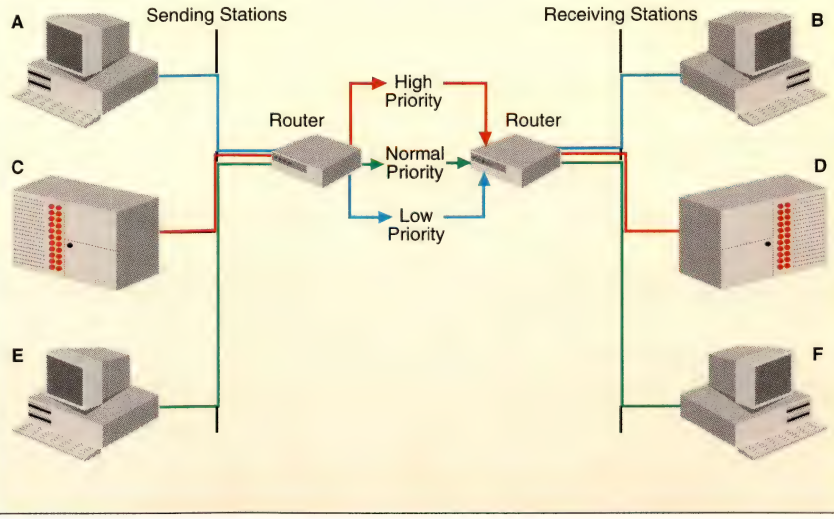

Which publication is read by more communications professionals?

ATUG - The Exhibition and Conference of the Australian Telecommunications Users Group is unquestionably the major communications event in Australia. Visitors to the annual event are widely considered to be the region's top communications professionals, responsible for the lion's share of its networking and telecommunications purchasing.

Upon registering at ATUG'93 visitors were requested to indicate the communications/networking publications they regularly read. Judging from the **4,690*** completed questionnaires *Australian Communications* is clearly the choice of communications professionals

" I regularly read...."

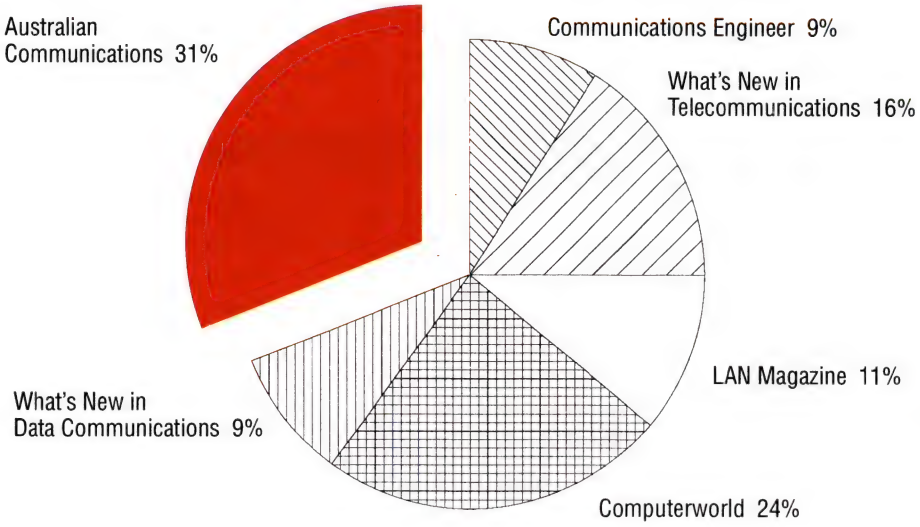

Data supplied by Riddell Exhibition Promotions, organisers of ATUG'93

* Does not include conference delegates, pre-registrations, or visitors with completed coupons from advertisements in *The Australian* and *Australian Communications.*

AUSTRALIAN communications
The Data and Telecommunications Management Magazine

International Advertising Representatives

Head Office
AUSTRALIA & NEW ZEALAND
Australian Communications
4/541 Kent Street
Sydney NSW 2000
Australia
Ph (02) 264 2200 Fx (02) 264 2244

USA & CANADA
M&T International Marketing
Miller Freeman Inc
411 Borel Drive Avenue
San Mateo CA 94402
Ph (415) 358 9500 Fx (415) 358 9739

EUROPE
A-Z International Sales
70 Chalk Farm Road
London NW1 8AN
United Kingdom
Ph 71 284 3171 Fx 71 284 3174

JAPAN
Media Sales Japan Inc
Tamuracho Building
3-3-14 Shimbashi
Minato-ku Tokyo
Japan
Ph 3 3504 1925 Fx 3 3595 1709

TAIWAN
Acer TWP Corp
4F 19-1 Lane 231
Fu Hsing N. Road
Taipei 104
Taiwan
Ph 2 713 6959 Fx 2 715 1950

KOREA
Jes Media International
5th Floor Banpo Building
40-7 Chamwon-Dong
Seocho-Gu Seoul 137-030
Korea
Ph 2 545 8001/2 Fx 2 549 8861

HONG KONG
Acer TWP Corp HK
Unit 2, 5th Floor
82-84 To Kwa Wan Road
Kowloon
Hong Kong
Ph 764 3857 Fx 764 3857

SINGAPORE
Eastern Publishing Associates
1123 Serangoon Road #03-01
Singapore 1232
Ph 296 6166 Fx 298 7551

ISRAEL
Nora Lifton
Kfar Saba 44444
Israel
Ph 52 970065 Fx 52 953194

Vendor Profiles

Below is a short description of each router tested by the Lab, and the approach taken by each vendor. All of the products tested are available in Australia now.

ADVANCED COMPUTER COMMUNICATIONS

Advanced Computer Communications' Express Queuing delivered the most sophisticated protocol prioritisation of any vendor, making it possible to share bandwidth and set explicit priorities. The feature worked in some of the Lab's tests, although it also produced some strange results. At this point the Lab doesn't know if these anomalies are caused by flaws in the product or problems with the tests.

Overall, the ACS 4200 is a straightforward bridge/router that handles the most common protocols. It can distinguish between different bridged protocols and was thus tested with SNA, NetBIOS, and IP.

Prioritisation is activated by issuing commands from an ASCII terminal. The interface is fairly straightforward, though the syntax is cryptic. The Service Access Point (SAP) value must be entered manually, and some commands are so long that they don't fit on a single line. Since commands can't wrap to the next line, each word had to be abbreviated. Since ACC does not list these abbreviations in its documentation, it took some trial and error to come up with a workable solution.

ACC's Express Queuing determines how many sessions are using the wide area link and then allocates bandwidth to each conversation. It forwards one packet from one conversation, another from a second conversation, and so on, thus making sure that sessions that would be blocked with a first-in, first-out algorithm receive a fair share of the bandwidth no matter how quickly packets arrive at the ACS 4200.

Express Queuing is the ACS 4200's default mode of operation. It can be turned off in favour of first-in, first-out routing, but the Lab sees little reason to do so.

Express Queuing also allows three levels of priority to be set — high, normal, and low. Normal means no priority. If a protocol is designated as high priority, the ACS 4200 transmits four packets for every two normal-priority packets (or four high-priority packets for every one low-priority packet).

It was easy for the Lab to get a good sense of how effective Express Queuing is. With no prioritisation set, SNA received 72% of the available bandwidth; NetBIOS, 1%; IP, 31%. (As discussed above, optimisation makes it appear that more than 100% of the 56Kbps pipe is being used). Express Queuing evened things out greatly: SNA received 59%; NetBIOS, 14%; IP, 27%. The Lab isn't sure why the bandwidth wasn't divided in thirds. As noted, one of the variables that can affect performance is the way a protocol reacts to congestion on the network.

When SNA was designated as high-priority traffic, it ended up with 82% of the bandwidth. NetBIOS was virtually unchanged at 16%. IP throughput dropped to 7%. The Lab cannot explain this anomaly. When tested with small IP packets, SNA received 72% of the available bandwidth; NetBIOS, 17%; IP, 20%. This seems to indicate that the rate at which packets are fed to the ACS 4200 affects Express Queuing and the way it prioritises packets.

The Lab tried assigning several other priorities. SNA was kept high, while NetBIOS and IP were set low. The Lab expected that NetBIOS and IP would get less throughput. Instead, NetBIOS remained nearly the same; IP went up to 29%; SNA dropped to 48%. The extra processing associated with filtering three protocols may account for the SNA drop.

ACC is one of only two vendors to offer data compression. This can greatly boost throughput across the board.

CISCO SYSTEMS

Cisco, the leading router vendor, has what looks to be a fairly straightforward priority feature based on precedence queuing. Unfortunately, the baseline measurements the Lab made revealed wild fluctuations in performance. Further, the Lab could detect little difference with prioritisation enabled or disabled. Thus, the tests are inconclusive.

The Cisco 4000 is a relatively new midrange platform that supports a maximum of four WAN ports. The 4000 continues the Cisco tradition of support for an extensive range of protocols. Prioritisation can be configured fairly easily from an ASCII terminal. The 4000 can distinguish between routed protocols but not between bridged protocols. The company says this shortcoming will be addressed in a future software release. The Lab tested the Cisco 4000 using SNA, IPX, and IP.

Priorities can be set at high, medium, low, and normal. The last setting falls between medium and low, and is the equivalent of no prioritisation. Queue size and latency can be adjusted to keep low-priority traffic from being blocked.

With no prioritisation, SNA received almost no bandwidth; IP ended up with 11%; IPX, 72%. When 64-byte IP packets were used, SNA bandwidth jumped to 79%; IP received 21%; IPX, 2%. The Lab is at a loss to explain these results for baseline performance.

In tests with SNA set to high priority and IP and IPX set to low, the same inexplicable allocation was noted. With 2-Kbyte packets, SNA received 9% of bandwidth; IP received 1%; IPX, 93%. With 64-Kbyte IP packets, SNA took 71%; IP, 29%; IPX, 44%. Similar results were seen with SNA priority set to high, while IPX was set to medium and IP to low. When 2-Kbyte packets were used, IPX's share of the bandwidth was 19%, SNA dropped to 9%, and IP ended up with 91%. With 64-byte IP packets, SNA received 71%; IP, 29%; IPX, 44%.

Given these erratic results, it's impossible for the Lab to make a definitive statement about Cisco's prioritisation. All of the aforementioned variables may be skewing the Lab's findings. In addition, test results may be affected by the Cisco 4000's apparent inability to handle 2-Kbyte IPX packets. The Sniffer LAN Analyser showed that the Cisco 4000 was actually fragmenting the packets, even though it was set to pass packets as large as 2.5 Kbytes.

Cisco also offers a way to prioritise SNA traffic from FEPs (front-end processors) and terminals. This feature uses SNA's own prioritisation settings.

CROSSCOMM

Thanks to its graphical user interface, Crosscomm has made sure that its prioritisation is easier to use than any other vendor's. Unfortunately, since the Crosscomm router allocates most of the available bandwidth to SNA traffic — even without prioritisation enabled — it's difficult to draw any conclusions about the feature's effectiveness. Configuring prioritisation is a matter of simply clicking items on a menu. It's not necessary to look up SAP addresses.

The Lab tested the router in its Protocol Independent Routing (PIR) mode. PIR is Crosscomm's proprietary Shortest Path First (SPF) routing algorithm adapted for use with almost every popular protocol.

Priority is based on precedence. There are three user-definable priorities. In tests with no priority set and using 2-Kbyte frames, SNA throughput was 96% of the 56Kbps line; IPX received 20%; NetBIOS, 18%. (Again, header stripping and other optimisation techniques boost throughput beyond 100% of available bandwidth). When 64-byte packets were used, total throughput dropped significantly, as did SNA throughput. SNA throughput came in at 95%; IPX, 1%; NetBIOS, 9%.

The Lab cannot explain why SNA received such a large portion of the bandwidth. It may be that PIR is better suited

Vendor Profiles (Continued)

to bridged protocols. Given Crosscomm's Token Ring background and IBM slant, it may simply be that its router is more adept with SNA than other protocols.

With SNA prioritisation set to high, SNA throughput went up only slightly. IPX and NetBIOS remained virtually unchanged. With 2-Kbyte frames and SNA set to high priority, SNA took 98%; IPX, 20%; NetBIOS, 9%. With 64-byte IPX frames, SNA ended up with 67%; IPX, 1%, NetBIOS, 7%.

IBM

IBM's prioritisation scheme is included in the Data Link Switching (DLS) capability of its 6611 router. It works only with NetBIOS and SNA traffic. When the 6611 was working with 2-Kbyte packets, the Lab was unable to detect any significant difference in performance with Data Link Switching on or off.

Configuration is via a Windows-based program, which first creates a configuration file that is then read by the 6611. All other vendors offer real-time configuration via Windows or ASCII terminals. The program is relatively easy to use. The 6611 also can be configured in real-time using an ASCII terminal, but the terminal interface is much harder to work with.

DLS is the way in which the 6611 optimises protocols commonly found in IBM computing environments. It's intended as an alternative to the 6611's source route bridging mode. With DLS, SNA and NetBIOS traffic are guaranteed at least 50% of available bandwidth. In this way, they are given priority over protocols that might block them if source route bridging and first-in, first-out algorithms were being used. DLS is an all-or-nothing proposition: it's either enabled or disabled. It's not possible to set priorities for traffic coming from different SNA or NetBIOS applications. All SNA and NetBIOS traffic is lumped together.

The 6611 was tested using SNA, IPX, and IP. Tests were first run using source route bridging (no prioritisation) and then using DLS. These were the only tests conducted by the Lab, since the 6611 cannot prioritise IPX or IP. With source route bridging and 2-Kbyte packets, SNA ended up with 18% of the bandwidth and IP received 102%; no IPX packets were passed. The 6611 cannot handle IPX packets longer than 1,500-Kbytes.

The same results were observed with DLS and 2-Kbyte frames. SNA and IP throughput were unchanged; IPX once again did not get through. With 64-byte IP packets using source routing, SNA received 29% of the bandwidth; IP, 130%; IPX, 1%. With DLS enabled, SNA was granted 41%; IP, 120%; IPX, 4%.

PROTEON

Proteon's priority scheme, like ACC's, melds bandwidth sharing and precedence. Proteon also makes it possible to set priorities for different types of IP traffic — something that no other vendor allows.

The Lab tested Proteon's CNX 500 (the brand new CNX 600 was not available at the time), which comes with four WAN ports. Prioritisation is configured via a command-line interface on an ASCII terminal. There are only a few commands, but they are very difficult to interpret.

Network managers can determine the amount of bandwidth dedicated to each protocol. Bandwidth not being used by one protocol can be borrowed by another.

Four priorities also can be assigned to different types of traffic that use the same protocol — urgent, high, normal, and low. Thus, 60% of the available bandwidth could be assigned to IP, with priorities established for Rlogin (TCP/IP login function), SDLC relay, IP tunnelling (IP encapsulation), and Telnet traffic. The Lab did not test this advanced feature.

Proteon is able to distinguish between different bridged protocols. The Lab tested the CNX 500 using SNA, IP, and NetBIOS. With no bandwidth sharing specified and using 2-Kbyte frames, SNA ended up with less than 1%; IP got 48%; NetBIOS, 65%. With 64-byte IP packets, SNA got 12%; IP, 106%; NetBIOS, 12%.

Bandwidth sharing worked well when it was activated. Using 2-Kbyte packets and awarding SNA 80% of available bandwidth, SNA throughput climbed to 68%; IP dropped to 15%; NetBIOS rose to 16%. With 64-byte IP packets, SNA took 29%; IP, 88%; NetBIOS, 7%.

RAD NETWORK DEVICES (RND)

RAD Network Devices offers fairly rudimentary prioritisation. Unfortunately, the Lab was unable to establish a performance baseline, since the RND router could not handle the large amount of data present on the test bed. Thus, the Lab cannot draw any conclusions about the effectiveness of RND prioritisation.

The Lab tested RND's Remote Token Ring Bridge (RTB-40). The RTB-40 is similar to Crosscomm's ILAN in that it uses a proprietary algorithm (RND's version of Shortest Path First) for all protocols (including bridged protocols).

RND's prioritisation was the most difficult to configure of any vendor tested. Users are forced to convert hexadecimal SAP values to binary; there's no excuse for putting network managers through this sort of trouble.

RND's prioritisation is based on precedence. Only one priority can be specified — high. (The low setting is the same as no priority.) The RTB-40 can distinguish between different bridged protocols. The Lab tested the unit with SNA, NetBIOS, and IP.

With no priority set and using 2-Kbyte packets, the RTB-40 was unable to handle the massive amount of data the test bed sent to it. The SNA and NetBIOS sessions crashed and IP was awarded 100% of the bandwidth. The same thing happened with 64-byte packets. When SNA was designated as the high-priority protocol and using 64-byte IP packets, SNA received 44% of bandwidth; IP, 36%. No NetBIOS data got through. With 2-Kbyte packets, SNA received 2%; IP, 91%. Again, no NetBIOS traffic got through.

WELLFLEET COMMUNICATIONS

Wellfleet's priority option is based on precedence. The Lab tested the Backbone Link Node, a unit that delivers very high performance and excellent fault tolerance.

Configuration with the company's Site Manager graphical interface was tough. When it comes to prioritisation, it's a nightmare of nested pull-down menus and pop-up windows that appear in response to selections. Worse, it was often impossible to determine which questions in which windows had to be answered first. If the wrong choice was made, necessary windows disappeared.

Prioritisation is based on queues. The high-priority queue is always emptied before lower priority queues. Three priorities can be set: high, medium, and low. Medium is the same as no prioritisation. To keep high-priority traffic from blocking the line, queue size and latency can be programmed.

Wellfleet's routers cannot distinguish between bridged protocols. Thus, tests were made using SNA, IP, and IPX.

With no priority set and using 64-byte packets, SNA received 67% of available bandwidth; IP, 22%; IPX, 2%. When 2-Kbyte packets were used, SNA throughput plummeted to 7% IP jumped to 95%; IPX climbed slightly to 4%. The Lab has no explanation as to why SNA throughput was so affected by IP packet size.

SNA was then set to high priority and IP and IPX set to low. When 2-Kbyte packets were used, SNA received 67% of the bandwidth; IP, 32%; IPX, less than 1%. With 64-byte IP packets, SNA jumped to 95% and IP plunged to 1%. No IPX traffic made it through.

Considering FDDI?

Test FDDI (And Everything Else!) With The DA-30 Analyzer.

FDDI doesn't exist in a vacuum–it is used in today's mixed topology networks. Don't invest in test equipment that can handle only part of your internetwork.

The DA-30 gives you the best FDDI analysis value and the modularity to add Token Ring, Ethernet, V.35, X.21, T1, RS-232, V.36/RS-449, SMDS and Frame Relay at 2 Mbps. Configure it to your network today and then expand it for the future.

For FDDI testing you'll be able to simultaneously transmit and receive traffic at the speeds this technology demands. The DA-30 collects information about network events such as frame counts, utilization, frame size and ring mapping. In either dual attachment mode or single attachment mode, you'll be able to decode MAC, SMT, LLC and upper layers running over FDDI.

The unique DA-30 dual analyzer tests mission-critical interconnect devices like bridges and routers at loads ranging up to 446,525 fps.

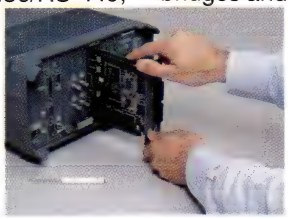

Before finalizing your FDDI decision, learn about the analyzer that tests your entire internetwork by calling

Melbourne: (03) 690 6700
Sydney: (02) 437 5011

E/1192/D3/010/4C/AUS

Wandel & Goltermann Pty. Ltd.,
PO Box 419
World Trade Centre, Victoria 3005

Wandel & Goltermann
Electronic Measurement Technology

Figure 3: Share and Share Alike

Dynamic sharing algorithms enable routers to keep tabs on traffic flow and distribute bandwidth according to need. Since every application is regularly serviced, there's no chance that a session will be blocked by congestion or terminated accidentally.

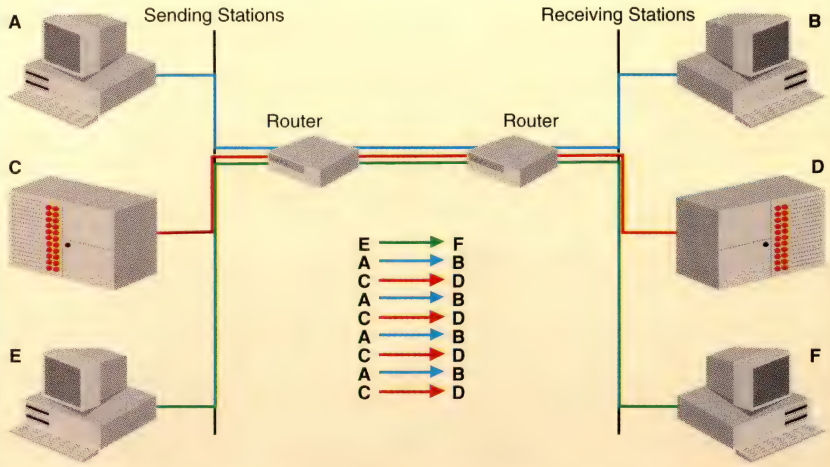

From page 75

tocols like NetBIOS. When they prioritise SNA, they're also treating NetBIOS as a high-priority protocol, thus defeating the purpose of the algorithm.

None of the products tested had any trouble telling the difference between various routed protocols. By definition, routed protocols can be identified readily by their Layer 3 addresses.

Whether a product makes it easy or difficult to assign priorities also is an important issue. Obviously, calling up a graphical interface, clicking on an icon, and answering a few questions is a lot simpler than struggling with bit patterns from a command-line interface. Crosscomm is the only vendor to supply an easy-to-use graphical interface. Wellfleet's interface — though graphical — proved to be unnecessarily complicated. All the other vendors resort to character-based interfaces. With most of these, network managers have to know and enter the value of the SAP (usually listed in the documentation). Once the SAP is entered, all configuration programs but RND's automatically create a filter that looks for the proper value. With RND, users have to convert the SAP value to binary and enter a bit mask that the RND filter can use to scan data. Forcing network managers to do hexadecimal-to-binary conversion seems a little out of place in this day and age.

The Test Bed

The Lab set up two 16Mbps Token Ring LANs connected with the router pairs under test. CSU/DSUs (channel/data service units) were used to create a US-style 56Kbps wide area link between routers. Clients and servers were placed on each Token Ring to generate multiple traffic streams, and an Expert Sniffer LAN Analyser from Network General measured the traffic forwarded by the routers.

A DA-30 LAN analyser from Wandel & Goltermann generated IP traffic. To create IPX (Internet Packet Exchange) traffic, a file was copied between a DOS workstation on one LAN and a server running NetWare 3.11 from Novell on the other. To generate NetBIOS traffic, a second file was copied between a DOS workstation and a server running IBM's LAN Server network operating system. SNA traffic was generated using one of the LU 6.2 Bader Benchmark programs, which consists of a client and a server application that move data via SNA file transfers.

The effectiveness of each router's prioritisation was tested by simultaneously sending three different protocol streams to the box under test. If the router could distinguish between bridged protocols, SNA, NetBIOS, and IP were used. If the router grouped bridged protocols together (as did gear from Cisco and Wellfleet), then IPX was substituted for NetBIOS. The Lab verified that all the routers were operating properly before the prioritisation tests started.

In order to establish a baseline for performance, the routers were first sent all three protocols without any priorities set. Next, SNA was prioritised. If different levels of priority could be established, that feature also was evaluated.

The packets for all tests were set at approximately 2 kilobytes. This length was chosen since it's representative of the packets typically found on corporate networks. To ensure that prioritisation wasn't implemented simply by transmitting more high-priority packets than low-priority packets, each router also was tested using 64-byte IP packets along with 2-Kbyte packets for the other protocols. If a vendor were using this simple approach to prioritisation, more small packets would be passed onto the network (since they out-number larger packets by such a large degree).

The Sniffer LAN Analyser was set to capture all data. Throughput for each protocol was determined by examining a 10-second slice of the capture. If the client and server completed the exchange in less than 10 seconds, the slice was smaller.

Test Results

As might be expected with the industry's first look at protocol prioritisation, it proved very difficult to quantify precisely the effectiveness of this new feature. For one thing, no standard test program exists that can simulate the sort of congestion that prioritisation algorithms are intended to handle. This forces applications to be used to generate data — a less than ideal approach that opens the door to unforeseen variables.

Differences in the behaviour of protocols can also skew test results. SNA, for instance, will throttle back if it detects that traffic is not getting through; IP, on the other hand, just keeps blasting away (especially when a frame generator like the DA-30 is being used). If SNA backs off, it will likely end up with *less* bandwidth — even if it has been designated as high-priority traffic.

Whether or not a protocol can be routed may also affect test results. Bridging requires less processing than routing. Thus, throughput for bridged protocols (SNA and NetBIOS) may be higher than for routed protocols, with or without prioritisation.

Differences in the way a router handles a particular protocol can also affect test results. During the course of its tests, the Lab noticed that it was getting more data through the 56Kbps WAN line than expected. Since it sometimes appeared as if routers were using more than 100% of the available bandwidth, the Lab first thought that some form of data compression was being used. It turned out that the test results can be explained by other techniques used to optimise throughput. Some transmitting routers, for example, actually strip the headers and trailers from specific protocols, which means that all the bandwidth in the pipe is being used to carry data. The receiving routers attach new headers and trailers. These router pairs can actually strip and attach so quickly that throughput ends up well above expected levels.

Kevin Tolly is Director of the Data Comm Test Lab and President of Interlab (Sea Girt, New Jersey). Eric Hindin is Senior Technology Editor for Data Communications magazine. Andy Hacker, Wayne Schiller and John Curtis of InterLab handled the testing and research for this Data Comm Lab Test.

To network your branch offices you need one of two things.

The resources.

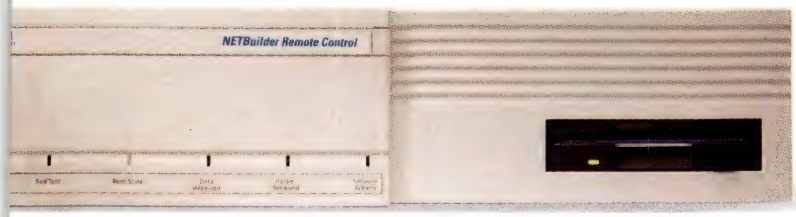

Or the brains.

Introducing revolutionary Boundary Routing from 3Com.

A sales office in Canberra wants data from headquarters. A branch office in Cairns wants e-mail.

Before, it took an army of network administrators. Now, it's simple.

Introducing Boundary Routing™ System Architecture from 3Com. A breakthrough software technology and a new remote router family that offer you the most cost-effective and simple way to add branch offices to your corporate network.

Now installing and maintaining your branch office networks has never been easier. Protocol changes? No problem. You can reconfigure your branch office networks by changing only the central routers. No box-by-box changes. And no need to touch the remote sites.

And that's just the start of the advantages.

To learn more, call **(02) 959 3020**, ext. 22 or contact your nearest 3Com distributor, and find out how smart connecting branch offices to your network can be.

3Com®

Networks That Go the Distance™

Tech Pacific Australia
(02) 697 8666

Tech Pacific New Zealand
(09) 415 8300

MM Data Networks Australia
(02) 980 6922

BICC Communications New Zealand
(09) 358 0322

i.t. connXions Australia
(02) 415 0555

Why not get the WHOLE story on networking?

Whether it's FDDI, OSI, ATM, ISDN, SNA, Frame Relay, Fastpac, Token Ring or 10Base-T, you need to be on top of the rapid changes now taking place in the communications industry.

In the pages of Australian Communications, every month the region's top communications journalists, leading industry figures, academics, professionals and industry experts bring you the latest news and analysis on the technical, political and management issues facing this dynamic industry.

The time to subscribe is NOW! with savings up to 20% off the newsstand price.

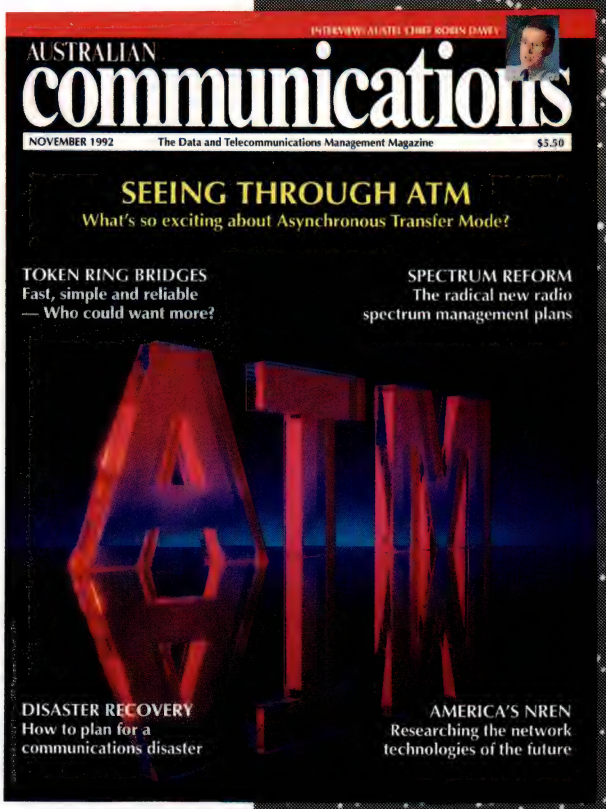

SUBSCRIBE NOW

☐ 22 issues, two years, for $96 (save 20%)
☐ 11 issues, one year, for $54 (save 10%)

I enclose my cheque for AUS$.....

☐ Bill my company Authorising signature.....

Charge my: ☐ Amex ☐ Diner's ☐ Visa ☐ Bankcard ☐ Mastercard

No ..... Expires.....

Name .....

Company .....

Address .....

..... Postcode .....

Signature .....

Post to: Australian Communications, Level 4, 541 Kent Street, Sydney NSW 2000
OR: Fax your application to us on (02) 264 2244

Will GSM and D-AMPS Give Way to the CDMA Push?

In the second part of an extended look at mobile communications technology, Stewart Fist investigates CDMA, the technology which may well take the place of the TDMA-based GSM and D-AMPS.

The problems that are being experienced around the world with GSM and Digital AMPS (see 'GSM — Better Late Than Never?,' the first installment of this article in the June 1993 edition of *Australian Communications*), are to do with time division multiple access (TDMA) techniques. In conversation with Americans, you'll find that they don't understand the term Digital AMPS (D-AMPS) at all — they use the generic 'TDMA' to mean their specific version which was chosen a couple of years ago by the Telephone Industry Association, as *the* US standard.

To be strictly correct, GSM and D-AMPS base stations use TDM (time-division multiplexing) on the outbound signals, and TDMA on the inbound ones from the mobiles. It's a bit of a esoteric distinction, but it is important to distinguish the ease with which you can fit eight GSM (three in D-AMPS) base-to-mobile channels inside time-slots when you are generating all the frames yourself (TDM), and the critical problems of controlling eight individual mobiles, at varying distances from the base, each trying to advance and retard synchronisation on the return links. They are all required to transmit only short bursts of signals in such a way that they arrive back at the base and fit perfectly into the time-slot allocated for them — which is TDMA.

It is this fact that makes the GSM and D-AMPS 'air-interface' so fragile. It takes a finite time for the synchronisation signals from the base to reach distant mobiles, and the same time again on the return path. So the only way to calculate distance, and therefore the necessary synchronisation advancement needed by the more distant units, is for the base to constantly transmit synchronisation bits, which will be echoed back to base (together with the mobile's identification) so that the base can monitor the delay. It can then advise the mobile to advance its synchronisation by this amount to avoid data overflowing into the adjacent time slots.

Even when the system works well, problems arise when a mobile first initiates a call. At that time, it hasn't had a chance to realign its slot timing. For this reason, GSM has an 'access burst' which uses only 64 bits at the front of the slot, and leaves the remaining space free. So if it over-runs, it is not scattering alien data through channels belonging to adjacent users.

And it is the distance over which radio waves travel during this 'half-slot' period left unoccupied which sets the maximum size of a GSM cell. If cells are larger, then alien access-pulse data will start to intrude into the adjacent slot. Fortunately, if mobiles attempt to transmit from more than 35km, their signals are likely to be too low in level to have much effect.

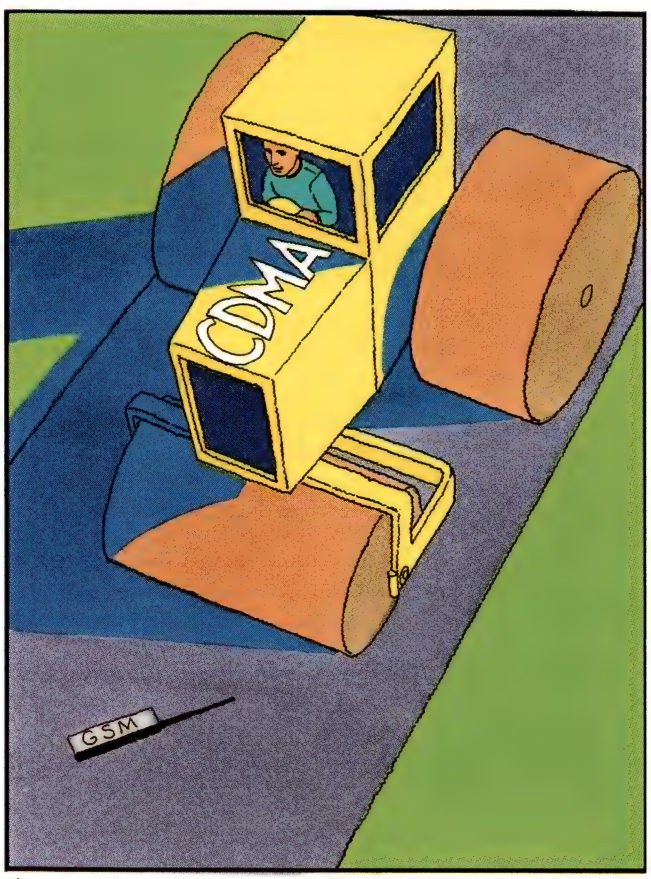

So with this cell size limitation, GSM wasn't the ideal technology for the wide-open spaces of the sunburnt country anyway. But, now this defect looks trivial. And it is solvable, anyway, by changes ('slot-stealing') due in the GSM Phase 2 supplementary standards in 1995-6.

What turned out to be the problem with all TDMA systems, however, wasn't this bit-intrusion or any other of the more esoteric issues which occupied so much of the development time on both

Can GSM's Problems be Solved?

The early tests all showed that GSM suffered from out-of-band interference, and more than expected spread of signals into neighbouring cells. This was drastically reducing the capacity of the system, and since the main thrust of the developments had been to achieve a 10-times increase in capacity-per-MHz of bandwidth, this was seen as a serious matter. They were only getting figures of 2x and 3x, even under test conditions.

In an attempt to boost capacity, DSI (Digital Speech Interpolation) techniques were introduced to reduce load on the channels by switching off the transmitter whenever a speaker paused for a moment, or was listening. So a more irregular burstiness of the transmitter was introduced to the tones being generated by normal frame-burst.

It appears now that the employment of this DSI technique will create even more problems with TDMA, since it introduces a randomness into the low-frequency component of the emissions. It may be possible to take DSI out of the sets, but only by reducing the capacity of the system even more.

It seems the only overall solution to these problems is to turn the handset and mobile transmission power down rather radically, and this means that cell-sizes will need to be drastically reduced in area, and, of course, in power.

In turn, this is going to produce problems. You can't have a 50 metre high antenna tower every half-kilometre or so throughout a city, so the smaller microcell type transmitter and lamp-post type units will need to be employed. But these will then have to be much closer to passing pedestrians.

If base station transmitter antennae are spread throughout a city at this height, it is likely that they will generate their own problems with the hearing-impaired passing by. I'm not too sure about this suggestion, because the bases transmit an almost constant stream of slot-signals, while the mobiles only transmit one-in-eight. It may be that the technical problems may be solved by using these small cells with low power and discarding DSI.

This won't solve the cost disadvantage, though. Ideally, what the carriers want to be able to do in the early stages of introducing any new cellular technology is to cover the area with the largest cells possible, then, as load comes on to the system and the funds start rolling in, they can progressively reduce cell sizes in the loaded areas, and increase capacity to keep pace with requirements.

Stewart Fist

D-AMPS and GSM. It was the basic physics of switching a transmitter on and off to generate the slot-burst.

Think of this in terms of the flicker effect of TV, which can send epileptics into a fit. It doesn't matter what channel you are watching, or what images the channel transmits — the flicker-effect arises from the 50 fields-per-second refresh-rate of the screen. In GSM, the on-off switching of the transmitter produces audio tones at 217Hz, and two harmonics above (434 and 651Hz), while American TDMA generates 100Hz, 200Hz and 300Hz (each channel has two slots in each frame).

A TDMA Flaw?

The important point is to understand how fundamental this low-frequency flaw is in TDMA systems. It will also exist in Japan's JDC standard, and in both half-rate coding D-AMPS and GSM cellular standards (GSM's is due in Phase 2), in the new DECT (Digital European Cordless Telephone) standard, and in the GSM-derivative DCS-1800 Personal Communications Network (PCN) which was supposed to revive UK electronics manufacturing.

The developers were so busy watching out for the high-frequency components that they missed the obvious low-frequency ones. Worldwide, the industry must be in absolute chaos at the present time. Gerry Flynn, the Chairman of the American TIA 45.5 subcommittee, almost sobbed down the phone: "It's a disaster worldwide — but for different reasons," he said.

The 'different reasons' he was referring to, are the associated problems of spurious emissions at the high-end of the radio spectrum. The on-off switching and the high bursts of R/F energy splatter interference into adjacent channels.

In Australia, they splatter into the existing analogue mobile channel, which is why we have more problems than the Europeans. But the Americans have even more of a crisis, because they have put Digital AMPS into reclaimed space in the existing AMPS spectrum — so they can't just abandon, say, a 5MHz segment of GSM-band to create separation, and move all transmissions to higher frequencies. We can.

Until now I have avoided writing up the health angle because there are ethical problems in beating-up fears which may not turn out to be based on fact. But I'm not going to avoid the discussion entirely.

Some medical researchers are now seriously looking at the long-term effects of microwave cellular (digital and analogue) in general. And digital's pulsed high-intensity, low-frequency signals are an almost perfect match for the ELF R/F specification that causes them the most concern. Don't forget these pulses come out of the antenna, so no amount of handset shielding is going to reduce them. Pulsed ELF frequency 'cooking' is the way in which chemists speed up some difficult chemical reactions, so digital, with both a pulsed and a higher-energy signal, must be a much more serious concern than analogue.

GSM, D-AMPS Under Threat?

So what does this mean for the future of digital telephony? First of all, it means that most of the long-touted advantages of both GSM and D-AMPS are fast evaporating. If these systems aren't already dead, they are seriously sick.

If, with their billions of dollars of R&D funds, GSM and D-AMPS researchers haven't managed to solve these problems by frantic work over the last two years, then it is unlikely that any new breakthrough is going to change the situation in the next 18 months, and by then CDMA is going to be a very serious rival.

So secondly, it means that Qualcomm's CDMA shares should shortly be skyrocketing! — (for this commercial advice, I'm charging 10%).

The American RBOCs and cellular operators are already turning away from D-AMPS, and a few are planning to use the Narrowband-AMPS (a 3x capacity, analogue system with digital controls) as an interim measure until CDMA can be introduced. PacTel has announced that it will use CDMA in Los Angeles, San Diego, Atlanta and Sacramento, beginning network installation later this year and commercial operations in late 1994. PacTel has eight of the top 25 US cellular markets, and it has placed an initial order for 30,000 dual-mode CDMA/AMPS handsets with Oki.

US West NewVector is planning an even earlier rollout in Seattle and some other unspecified locations. It has placed an order with Qualcomm for 36,000 dual-mode CDMA handsets, and it is already rolling out Narrowband-AMPS (N-AMPS) as an interim measure. Bell Atlantic also announced recently that it was abandoning TDMA in favour of CDMA (it services 15 States) beginning in Connecticut, Rhode Island and Massachusetts early 1994, and the Koreans (ever anxious to beat the Japanese to the punch) have also decided to go with CDMA on a national basis.

D-AMPS, under the original FCC ruling, was supposed to provide a 10-times capacity increase over analogue AMPS (taking into account the cell-reuse pattern, etc.) but it ended up with a 3:1 or 4:1 claim. In practice, I am told that many of the trials have shown that, under load, the systems can only provide 2:1 capacity gains, and even then the overall voice quality and ac-

ceptability (calls dropping, etc.) is inferior to AMPS. This is quite apart from their spectrum-splatter problems.

N-AMPS is currently better and cheaper, and it provides a 3:1 capacity gain without making much modification to the base stations. There are both N-AMPS and D-AMPS dual-mode (analogue AMPS also) handsets on sale in America, but only US West has N-AMPS up and running in the top-end of the analogue AMPS band.

They have taken back a chunk of spectrum, inserted N-AMPS to get its three-times capacity gain, and then are using this release of pressure on the rest of the band to allow them to change over to CDMA. Progressively, they will slip their customers over to N-AMPS and CDMA — getting 3x and 10x gains (respectively) out of these technologies, while continuing to support the remaining AMPS users.

GSM also started out with highly optimistic claims of capacity increases. But Hong Kong trials reduced expectations somewhat when they turned up figures of 1.5:1. Mind you, Hong Kong is the most difficult city in the world because of its constant canyons of concrete. However, some European trials have only shown 2:1 capacity increases, with a N=9 and N=12 cell reuse pattern. And, as Tim Harrabin of (UK) Cambridge consultant's Analysys says, "in these circumstances, why would you bother with GSM?"

The anomaly here is with Germany (and to a lesser degree Portugal). Germany supposedly has 400,000 GSM terminals in operation, and about 800 base stations. And there's no coincidence that both Germany and Portugal use an ancient analogue technology which operates in the 450MHz band. This gives them separation enough to avoid the splatter effect, although they now have broadcasters and electronic equipment manufacturers worried about interference in other areas.

Even densely-populated London has no capacity problems with the current Extended TACS analogue cellular system, and this fact certainly contributes to the reasons why the Brits aren't keen to roll out GSM — it's not just an altruistic concern for hearing-aid wearers.

The CDMA Way

But where most of the press (and almost all of the commentators) get it wrong, is in assuming that interference effects are a 'digital mobile' problem. You can't use the generic term here, it is a 'time-division' (TDMA) problem. CDMA doesn't pulse its transmitter, so it doesn't produce low-frequency interference, and since it uses much less power and spreads that power across a 1.25MHz-wide band, it doesn't produce much high-end R/F interference either. (For a full description of Qualcomm's CDMA technology see 'GSM — Why are We Rushing to Embrace it?' in *Australian Communications*, February 1992.)

Incidentally, you may hear the term 'wideband' being used in relation to CDMA, and it seems to be a generic term for any system which spreads its signals over a megahertz or so of spectrum, rather than using a narrow band like FM (analogue AMPS) or TDMA type systems. The new US TIA standard for CDMA 'wideband' is IS-95.

To insert CDMA into an existing analogue band they take back 1.8MHz, even though the wideband signal is only 1.25-MHz — which, if my maths are right, means that they are leaving 275kHz of guard space on either side (nine analogue channels). But this is to stop analogue mobiles from interfering with CDMA, not the other way around. CDMA actually operates below the normal noise threshold, so multiple CDMA channels can be butted up against each other, and are unlikely to produce interference with neighbouring bands.

Nor does CDMA seem to have suffered the same erosion of capacity that GSM and D-AMPS have demonstrated. PacTel, in its press release, says it expects 'over 10 times analogue, improved call quality, broader coverage and enhanced call privacy,' and US West says '10 to 15 times.'

And, when you look at the circuit components and understand the amount of messaging that is going on across the air-interface, you realise that these mobile phones (CDMA and TDMA) are extraordinarily complicated. We've reached the point where the complexity is too much for any single mind to grasp, which is possibly why a camel is coming out of the horse-committee. I think there's a 'philosophical' reason why these problems should be more difficult to handle in TDMA than in CDMA also.

The main problem with TDMA mobiles is with the fragility of the air-interface. Time-division access is an inherently difficult task to perform when you've got multiple users over varying distance, and with digital signals both reflected and direct, etc. Like hedgehogs having sex, the amazement is that they can do it at all. As a consequence, most of the complexity in the handsets and base stations was the result of

Spectrum Utilisation

A number of American carriers are closely monitoring US West's integration of CDMA and N-AMPS into the existing analogue AMPS spectrum. This company decided to deploy dual-mode N-AMPS/AMPS handsets as a way to extend the capacity of its band through Narrowband AMPS' 3-times capacity gain. It has been selling these dual-mode handsets since December 1992.

This approach has now freed up pressure on the spectrum enough to allow US West to take back 1.8MHz of bandwidth in the centre for its pilot CDMA deployment. The RBOC claims that tests show a 10 to 15-times capacity increase here, so it will then be easy at a later date to swap more AMPS bandwidth over to CDMA use.

Under FCC regulations, mobile carriers must continue to support all deployed standards well into the next century, but the capacity gain from CDMA should be more than enough to handle expected market penetration without needing to convert the whole band. Eventually US West expect to have 12.5MHz in use for AMPS (A, A' and A'') 10MHz for CDMA (B), and 2.5MHz for N-AMPS (B'). If necessary, the CDMA band could later extend up to the current N- AMPS band.

Stewart Fist

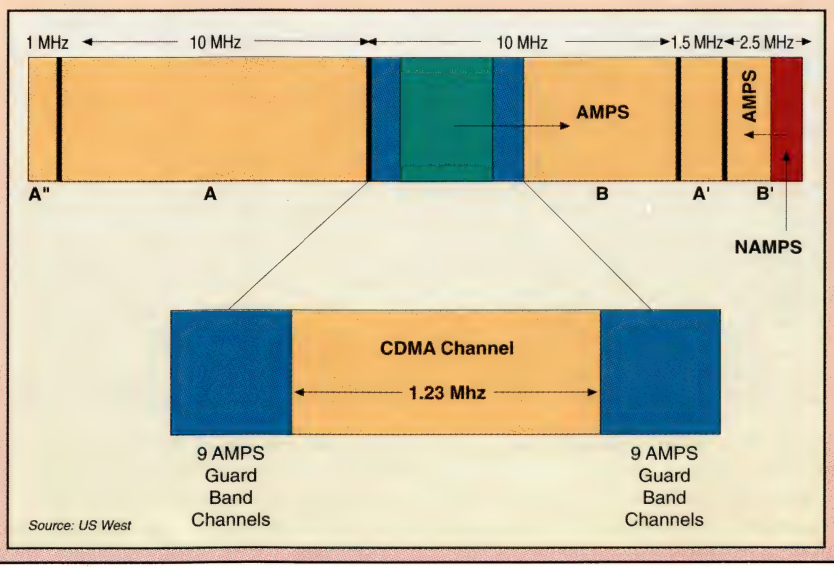

Smaller and Cheaper

No matter how good the technology, there's little likelihood of market penetration unless the equipment can be manufactured at an affordable price. And with all digital equipment, the primary decider of prices is the degree to which the various essential components can be integrated into the least number of chips.

California-based Qualcomm's expertise is in chip design. And the rapidity in which it is reducing the size and power-requirement of handsets through chip integration is quite exceptional — especially when you realise that these components are being designed to handle both the existing FM AMPS system, and the new CDMA.

Already in the first generation of dual-mode portables the chip-count for the CDMA components (CPU, CDMA signal processor, and vocoder) have been reduced to three, and the second generation chips are already at the foundry. Third generation handsets are already on the drawing board, and Qualcomm predicts handset costs at that time will drop to the $US100 to $US200 mark and become a mass consumer item for the emerging Personal Communications Services (PCS) networks.

Stewart Fist

Portable Phone Integration Path

	First Generation Dual Mode Portable	Second Generation Dual Mode Portable	Third Generation Dual Mode Portable
CPU	16 bit Microcontroller	16/8 bit Microcontroller	MSM 3.0
CDMA Signal Processing	MSM 1.0	MSM 2.0	
Vocoder	Mask ROM DSP		
FM Signal Processing Audio Circuits	Conventional ICs	Baseband Analogue ASIC Version 2.0	Baseband Analogue ASIC Version 3.0
Analogue-Dig Converter Dig-Analogue Converter	Baseband Analogue ASIC Version 1.0		
Baseband Filtering			
TX RF Modulator RX RF Demodulator		AGC ASIC TX ASIC RX ASIC	RF Front End MMIC
AGC (TX, RX)	Discrete		
RX Downconversion	Discrete		
TX Upconversion	Discrete		
Synthesizer	Hybrid (3)		
Power Amplifier	Standard AMPS Module	Custom MMIC	Custom MMIC

Source: Qualcomm

attempts to overcome these in-built air-interface instabilities.

With CDMA however, the air-interface is extremely robust and stable — it is shared among its many users, and the problems with reception are ones of noise, not timing. CDMA works below the background noise level, and relies on the signal being spread over a wide band or spectrum — so individual interference effects don't hit it hard.

CDMA's problems (and this was why so many experts ignored the technology) were to do with the need for massive computing power in the handsets to process the signals after they'd been received. The CDMA 'tuner' might be handling 20 different conversations on the same channel at the same time, and the selection of just the one that you wish to receive is a computer process. The computer is looking for a set pattern of bits (the 'code') and mathematically filters and processes only these. So CDMA has benefited from the increased number-crunching power of computers and the enormous advancements in digital signal processing — its problems were solvable by number-crunching.

And there's no finite 'call-blocking' here. It's not like GSM where, when eight users are occupying a time-divided channel, there are no more slots available. CDMA just keeps adding users, and the quality degrades 'gracefully' (their term, not mine) until those with the less important communications give up, terminate their conversations, and so remove one more source of the noise. It is self-regulating, without being limiting — that's the theory, anyway.

I should also point out, that the main perceived CDMA problem was transmission power level control. These are constantly monitored by the base station, and control signals are sent back to the mobile to turn power up or down so that all signals arriving at the base are within 1% of each other in terms of reception power. In fact, this is where Qualcomm has its main patents.

Certainly all the tests are showing that CDMA doubles battery life, and this is important in handhelds.

The emphasis on power control rather than timing control provides many advantages. Firstly it means that the battery isn't being wasted by mobiles close to the base station, and secondly it allows cellular systems to be almost infinitely scalable. So there's no reason why, in the future, you won't be using the same CDMA handset inside the office for a 30 metre link to the PABX; outside for Telepoint operations in the street; right through the handheld, transportable and vehicular cellular range; up to (perhaps!) direct satellite links when you are Back-o'-Bourke.

I'm certainly appreciating now why the line-up of carriers, terminal and base-station manufacturers at the recent Qualcomm/PacTel seminar in San Diego (February 1993) are the big names in American electronics.

Ericsson and Motorola, both of whom previously had been publicly scathing in their attacks on Qualcomm's CDMA proposals, have suddenly suffered a born-again conversion; AT&T and NorTel have gone into partnership (will wonders ever cease?); and the old reliable supporters, PacTel, Nynex, US West, Ameritech, Oki and Nokia were all in there cheering enthusiastically and giving papers of their trials.

One of the reasons I am reasonably confident that there won't be any unexpected disasters with CDMA is that Qualcomm has been very open in releasing details of the technology and trial reports of its development.

I hope for Nokia's sake that things work out well; its executives had virtually staked the company on the success of GSM, and so you can mark the European realisation of the extent of the GSM's interference problems by the date at which Nokia joined hands with Qualcomm.

Nokia was one of the early ones to begin looking elsewhere beyond GSM, just in case, while the other European manufacturers seem to have imitated the ostrich and

Pty Limited

PRESENTS

madam

The most advanced recorded voice announcer for Call queuing and ACD systems

Designed for PABX

Q's

008 809 359 Free Call
or Ph (02) 964 9555 Fax (02) 964 9988

YCOM PTY LTD
PO BOX 156
CROWS NEST
NSW 2065

Longer Battery Life

The usefulness of mobile handsets is generally limited by the life of the battery, so power consumption in both stand-by and talk-mode is important.

This is the graph of power consumption in a Bell Atlantic trial where measurements were being taken on a 'stereo' pair — one transmission path being AMPS, and the other CDMA. Both signals were between transmitters at the same cell site and mobiles in a single vehicle.

The FM AMPS mobile increases and decreases power in a limited number of steps of about 5dB each, while the fine power-adjustment exerted on the CDMA mobile by the base station moves it up and down rapidly over a much wider range. This record shows CDMA's transmission power is from 14 to 30dB below that of AMPS.

Generally, in practical tests, CDMA's talk-time battery life is extended about two- to three-times with the current first-generation handsets. With greater integration, and the introduction of 3V chip technology in the third-generation handsets, the standby time is expected to extend well beyond a working day.

Stewart Fist

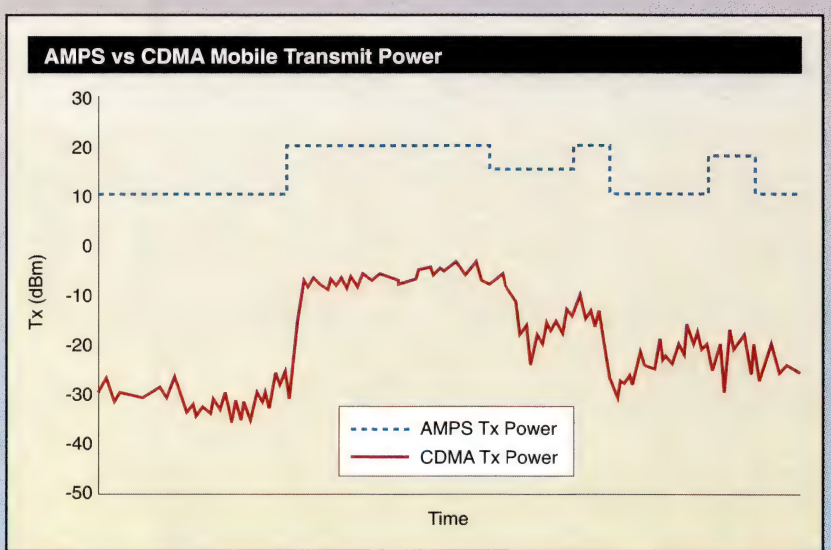

hoped the problems would go away. I'd be surprised if America doesn't hold the whip-hand now in the booming international mobiles market.

CDMA's Growing Support

From the start, Qualcomm has involved a number of carriers in all of the early trial phases, and AT&T, Motorola, PacTel, Ameritech, Nynex, MCI, GTE and US West have had staff almost full-time in San Diego running the two years of field trials. In this last year, San Diego must have looked like a constantly rotating Rotary convention: there's been so many people down there, running around in cars and trucks, recording voice, and tweaking electronics as they charge through the most difficult part of the city and surrounding canyon country.

The February conference was to present home-market field trials from Ameritech, US West, GTE, Bell Atlantic, PacTel and AT&T, and it marked the approval of the American Cellular Telecommunications Industry Association (CTIA) of the CDMA standard. The revised standard is only marginally different from Qualcomm's original proposal, and they tell me that they don't need to change any hardware. These are all software adjustments — so the changes won't delay any roll-out.

US West will be first off the mark with a pilot network running in Seattle later this year, and they'll boost this to a full commercial system when handsets become available mid- to late-1994. They may even hurry a bit faster now.

I've managed to get nine-hours of videotape showing the main speakers at the carrier's conference, with some edited highlights of the trials included. If this had been a Qualcomm-only propaganda tape, I'd still be cynical, but these trials were all being presented by the intended users.

Many of the recent trials included parallel transmissions of analogue (FM) AMPS and CDMA, or analogue AMPS and Digital AMPS, or D-AMPS with CDMA. They transmit the same voice through the same base station, at equivalent power-output, and receive them simultaneously with recordings made on a stereo DAT recorder and S-VHS videotape, so you are able to see and hear parts of the trials; the videotape lets you check both the external conditions and the signal quality measurements. The most interesting part is when the test vehicles are working out on the fringes of the cells or in 'vegetation tunnels' and in city streets with passing buses and trucks.

I don't think there's any doubt now that CDMA is intrinsically better on all counts than both analogue and Digital AMPS, and my viewing of the tapes suggests that D-AMPS is noticeably inferior to the old analogue system. I suspect D-AMPS has stopped dead in the water.

They also built a similar hand-mobile test unit and took it inside office blocks and various buildings to simulate someone walking around an office. Here CDMA also came out a pretty convincing winner. It doesn't seem to have the common 'Rayleigh' fading ('picket-fence') problem of TDMA systems, caused by standing waves of reflection from solid concrete walls and metal filing cabinets.

At 900MHz, the wavelength of the signal is only 30cms, and so a change in the differential path length between the direct and a reflected signal (or two reflected signals) of only 15cm can change their interference from constructive to destructive.

The other problem that engineers foresaw with CDMA was that the handsets would be complex and bulky, and possibly too expensive. It was assumed that most carriers would roll-out D-AMPS, and therefore CDMA may not reach the manufacturer's economies-of-scale needed to give it a chance of competing. Now, the economic argument has probably been reversed.

Qualcomm is fundamentally a chip-design company. It is also the recognised expert in CDMA-over-satellite technology, with its OmniTracks truck-location/data-communications system. And now Qualcomm and Loral have a joint venture to float a few dozen low-earth orbiting (LEO) satellites up in the sky in 1995-6 to compete with Motorola's Iridium voice-over-satellite network.

When I visited Qualcomm a year or so ago, the company had managed to condense down the main CDMA processing unit to three chips, and since then they've dropped the count to one. This leaves the R/F power units (which are reasonably conventional) as the most bulky item to fit inside the plastic housing, and the next generation of handsets will reduce the chip count even further. Senior Qualcomm executives believe that, in two to three years, cellular and PCS/PCN (Personal Communications Services/Networks) will merge into one, with pocket handsets selling in the $US100 to $US200 region.

Taken all together, CDMA is certainly the technology to watch.

Stewart Fist is a freelance journalist based in Lindfield, NSW.

Implementing Proactive LAN Management

If network managers are ever to feel comfortable entrusting vital corporate data to LANs, LAN management packages must go beyond mere troubleshooting and detect problems before they occur.

There are plenty of reasons to move corporate applications from mainframes to LANs. And there's one very good reason not to: LAN management still can't fully protect mission-critical data.

Users on mainframe networks rely on a wealth of behind-the-scenes services that keep applications running smoothly and efficiently: performance tuning, capacity planning, automatic alarm handling, software distribution, backup, and security. They take for granted that when new users come on-line, service won't be interrupted or performance degraded.

But such services are spotty at best on LANs, and all too often they're simply missing. LAN management typically means troubleshooting, and troubleshooting, by definition, implies that networks — and those who rely on them — are already in trouble. LAN management must do more: It has to be able to spot problems *before* they happen, instead of merely reporting when they do. It must, in a word, become proactive.

This isn't news to vendors of network monitors, protocol analysers, management systems, smart hubs, and simulation tools. They've started to add the underlying intelligence needed for proactive management. Most of the monitors now on the market, for instance, can track a range of network activity over time, giving net managers a way to create a profile of normal LAN traffic and spot changes that could adversely affect performance. Monitors that help network managers establish baselines and set alarm thresholds go a long way toward helping users anticipate hot spots and take steps to avoid trouble.

At the same time, monitor vendors are starting to make use of the wealth of data available with SNMP and other management protocols. Intelligence also is being added to other tools: Smart hubs are getting smarter, so that they can keep tabs on the multiple segments they control. And for the first time, PCs and workstations are being equipped with software that lets them report to popular management systems.

This is a good beginning. But if proactive LAN management is going to be more than a buzzword (and if LANs are ever going to be fully trusted with mission-critical data), monitors, analysers, and management platforms from various vendors must be capable of working together. Only in this way can proactive LAN management hope to match the suite of services found on mainframes. Key to this effort are guidelines emerging from groups like the Desktop Management Task Force and the Open Software Foundation, and standards bodies like the Internet Engineering Task Force (IETF).

What's the payoff? In the future, LAN management tools will establish a framework in which devices like smart hubs will monitor network activity and send information to distributed management platforms that will automatically issue trouble tickets, gauge operating costs, generate usage reports, and so forth (see the figure on page 98). These same platforms will feed information to simulation programs to help with capacity planning and configuration. Day-to-day activities will be carried out in the background, unseen.

A Sampling of Network Management Monitors

VENDOR	PRODUCT	DESCRIPTION	NETWORKS SUPPORTED	RMON SUPPORT
Armon Networking Dataplex (03) 210 3333	Armon OnSite	Proprietary multi-segment hardware or Sun/OS-based monitor. Unix-based management	Ethernet, Token Ring	Ethernet, Token Ring
Azure Technologies INS (02) 906 6335	LANpharoah LANpharoah Octopus	DOS software and adaptor cards for monitoring Ethernet and Token Ring 8-card expansion chassis and software	Ethernet, Token Ring	None
Hewlett-Packard 13 13 47	HP Openview Probe Manager and Lanprobe II	Proprietary hardware monitor and HP/UX or SunOS management software	Ethernet	Ethernet
Network Application Technology ADE Network Technology (03) 543 2677	Ethermeter and NMS/100E	DEC or Sun monitor with DOS management software	Ethernet	Ethernet
Network General Infotron (02) 417 7300	Distributed Sniffer System with Expert Analysis	Proprietary hardware monitor; DOS or Unix management software; expert system software	Ethernet, Token Ring, WAN	None
Novell Com Tech (02) 317 3088	LANtern and LANtern Services Manager	Proprietary hardware and Windows management software	Ethernet	None
Protools Ungermann-Bass (03) 696 2006	Foundation Manager and Cornerstone Agent	OS/2 monitor and management software	Ethernet, Token Ring, broadband, wireless	Ethernet, Token Ring
Telecommunications Techniques Vicom (03) 690 9399	NetLens 7000 Distributed Analysis System	Proprietary hardware monitor and Windows management software	Ethernet, Token Ring	Ethernet, Token Ring

Actively Proactive

Proactive LAN management means learning to make the most of today's tools. Most net management systems report on trends in network usage, number of errors, specific packet types, and other parameters over selected intervals, and also let users establish thresholds to flag aberrant network activity. HP Openview from Hewlett-Packard, Monet from Hughes LAN Systems, IBM's AIX Systemview NetView/6000, Racal Management System from Racal-Datacom, and Sunnet Manager from Sunconnect all offer these features.

Smart net monitors from Frontier Software Development and Protools offer similar capabilities, and these vendors and others are adding relational databases that make their products easier to use or to integrate with other management tools.

Another function, called baselining, is emerging as equally vital. By sampling activity over time and identifying normal performance through averages, means, and other statistical calculations, it is possible to establish a profile (or baseline) for any network. This baseline can then be used for a range of functions, from capacity planning to gauging the traffic levels in order to evaluate connectivity options.

Baselining may sound like another term for trend reporting, but it's broader in scope. It seeks to produce a statistically valid characterisation of normal network behaviour over an extended period rather than for a specific interval, taking into account varying levels of traffic at, say, different times of the day or on different days of the week.

A financial services institution, for instance, may consider it normal to leave sizeable chunks of LAN bandwidth idle in case of unexpected information dumps. Networks that mainly handle office applications or e-mail, in contrast, might run at higher levels of saturation. Normal activity for each would be defined in context rather than against a theoretical ideal, and baselines established. Network managers can then set meaningful thresholds to alert them to abnormal activity.

Some experts suggest that baselines should be checked and revised regularly to stay ahead of changing traffic patterns and keep them from becoming disruptive. Updated baselines, for instance, could reveal steady increases in LAN segment utilisation that indicate the need for further segmentation before performance suffers.

Not Enough RMON

Some baseline features are furnished by the SNMP Remote Network Monitoring Management Information Base (RMON MIB). RMON consists of several groups of MIB variables, each of which gathers performance data on a LAN segment through the first three layers of the ISO stack.

But the traffic statistics and historical data that RMON furnishes aren't all that's needed for sophisticated baselining. For instance, while most network monitors offer some means of averaging RMON data across multiple intervals (such as averaging performance criteria gathered every 10 minutes over eight hours), they can't statistically compare several days' worth of data. And they also can't assign different statistical weights to reflect expected deviations from 'normal' behaviour, such as increased network activity when a financial institution is calculating end-of-quarter reports.

Most network monitor vendors are working on the software needed to perform such tasks. In comparison, HP has equipped its Ethertwist hubs and bridges with what it calls an 'Embedded Advanced Sampling Environment,' firmware that continuously samples network activity, calculates normal parameters, and reports anomalies.

A market in baselining services also has formed, spearheaded by LAN integrators and service providers that supplement their own expertise with a combination of homegrown tools and commercial products.

Typical of these is EAC Connecting Point, a US LAN integrator that has come up with a baselining service based on Hewlett-Packard's Network Advisor protocol analyser. EAC will visit customer sites and spend up to two weeks sampling network traffic at selected times and running proprietary tests. These evaluations track net throughput for specific devices, protocols used, types of data, and frame sizes relative to throughput. Other measurements of network performance are also made. The company may ultimately develop a user kit for

DATABASE	INTEGRATED MANAGEMENT APPLICATIONS & PLATFORMS	PRICE
Flat file	HP OpenView, Sunnet Manager, Multiman	Manager Software, Software Probe, both P.O.A. Ethernet Probe available 09/93, Token Ring Probe available 12/93
Flat file	None	LANpharoah $9,500; LANpharoah Octopus P.O.A.
Flat file	HP OpenView, Sunnet Manager	$4,900
Flat file	DEC Polycentre SNMP Manager, Sunnet Manager	Ethermeter $5,400; NMS/100E $5,088; Kit containing Ethermeter and NMS/100E software $6,900
Flat file	Any SNMP management system	Monitor Server from $10,100; Analysis Server from $16,000; Sniffmaster Client from $9,500; Distributed Sniffer Starter Kit (2 x Servers, 1 x Client) from $45,000
Flat file (Btrieve)	None	LANtern $9,168; LANtern Services Manager $8,490
Flat file	Racal-Datacom Racal Management System, Ungermann-Bass Netdirector	Foundation Manager $20,600; Cornerstone Agent $2,900
Flat file	None	P.O.A.

the Network Advisor that includes a package of its tests.

Monitoring the Market

Despite their shortcomings, network monitors, or probes, are emerging as the linchpin of proactive management. Unlike management platforms, which keep tabs on specific devices, monitors continuously track packets actually crossing a network and gather details about the information they contain (such as size, number, and kinds of errors flagged; overall usage of a connection; number of hosts and their MAC addresses; and details about which hosts are sending information to specific devices).

Many monitors were derived from protocol analysers and can capture packets and decode protocols if necessary. But most simply filter and count packets. Monitors extend the range of management systems by gathering data at remote sites and sending it to the central management station. By reporting selectively or on request, they can save bandwidth (see table).

Most monitors use RMON MIB to gather data. Since the IETF has only approved a set of Ethernet variables for RMON, some monitors contain extensions for other networks. Some vendors, such as Protools and Telecommunications Techniques, have extended RMON to Token Ring. Some have gone further. Metrix Network Systems has extended RMON to FDDI networks, and Protools has adapted it for both wireless and broadband transmissions.

Vendors also are extending products in other ways, adding upper-layer information about the LAN protocols being used on a particular segment. Products that report on network traffic by protocol type allow users to determine the overall percentage of network traffic devoted to DECnet, Telnet, and other protocols. Frontier and Metrix also supply information about the performance of Network File System applications running on workstations from Sun. Both vendors gather RMON statistics by user-specified domains, and Metrix offers an application that allows users to view and set thresholds on traffic levels between segments.

Several monitor vendors that don't use RMON say they can match or exceed the functions of their competitors. For example, the Trakker product from Concord Communications uses SNMP but not RMON to gather data; its optional Trakreport package stores all data gathered by the monitor in an SQL relational database that can be imported into spreadsheets and other third-party applications.

Concord has long provided seven-layer monitoring and NFS support, as well as user-specified domains. Its Who Tool identifies workstation users and indicates how much bandwidth they're using, how many errors they're generating, and other details. Trakker also can identify the application-level protocols being used, as long as they're compatible with TCP or DECnet.

Novell asserts that the original RMON MIB was based on the data obtained by Novell's LANtern network monitor. Although Novell was active on the IETF committee that defined RMON, LANtern does not officially comply with the MIB.

At press time, however, Novell was developing a series of NLMs (NetWare Loadable Modules) that will support RMON and replace LANtern hardware with software. The latter communicates with the LANalyzer Services Manager, an application that will run in Novell's Netware Management System (NMS).

Getting Close to Hubs

Given their importance to proactive management, it's not surprising that monitors are being integrated into other management platforms. Some vendors offer their monitors as software for Unix hosts, and Frontier Software Development also has a DOS version. These software-only tools allow workstations or PCs to double as monitors.

Monitors also are moving into smart hubs, keeping watch over multiple LAN segments that would otherwise require their own probes. Metrix Network System's software serves as the basis of the Network Control Engine (NCE), a monitoring module that SynOptics Communications sells for its LattisNet hubs. Protools' monitors have been recast as modules for Ungermann Bass' Access/One. Hub maker Bytex ships an RMON monitor for its Series 7700 Intelligent Switching Hubs; the probe comes in Token Ring or Ethernet versions. DEC is rumoured to be readying an RMON agent for its DEChub 90.

Some manufacturers of smart hubs are putting RMON agents directly inside their gear, without adding monitoring modules. One of the first to do so was Star-Tek, whose hub business was recently bought by 3Com. And Frontier has announced an RMON agent expressly for smart hubs that supports Ethernet, Token Ring, and FDDI and can be managed via Windows, HP OpenView, and Sunnet Manager.

Just putting monitors inside hubs is not enough. Hubs should be able to reconfigure themselves automatically in response to information from monitors. To date, no hub on the market can do this.

Management by Committee

Monitors gather information and — to a limited degree — interpret it. But the applications that fully exploit the data pulled in from remote sites will run on management platforms such as HP Openview, Sunnet Manager and IBM's NetView/6000. Right now, many monitor and hub vendors sell integrated packages that let their tools be launched from the main menus or map displays of these platforms.

But for proactive management to become a reality, network managers must be able to share information among applications — to launch diagnostic programs

The Proactive Prognosis

In the future, proactive LAN management tools will gather data from across the network and automatically take actions in response to what they learn, reconfiguring hubs, issuing trouble tickets, simulating alternative configurations, and keeping net managers up to date about changing conditions and potential hazards.

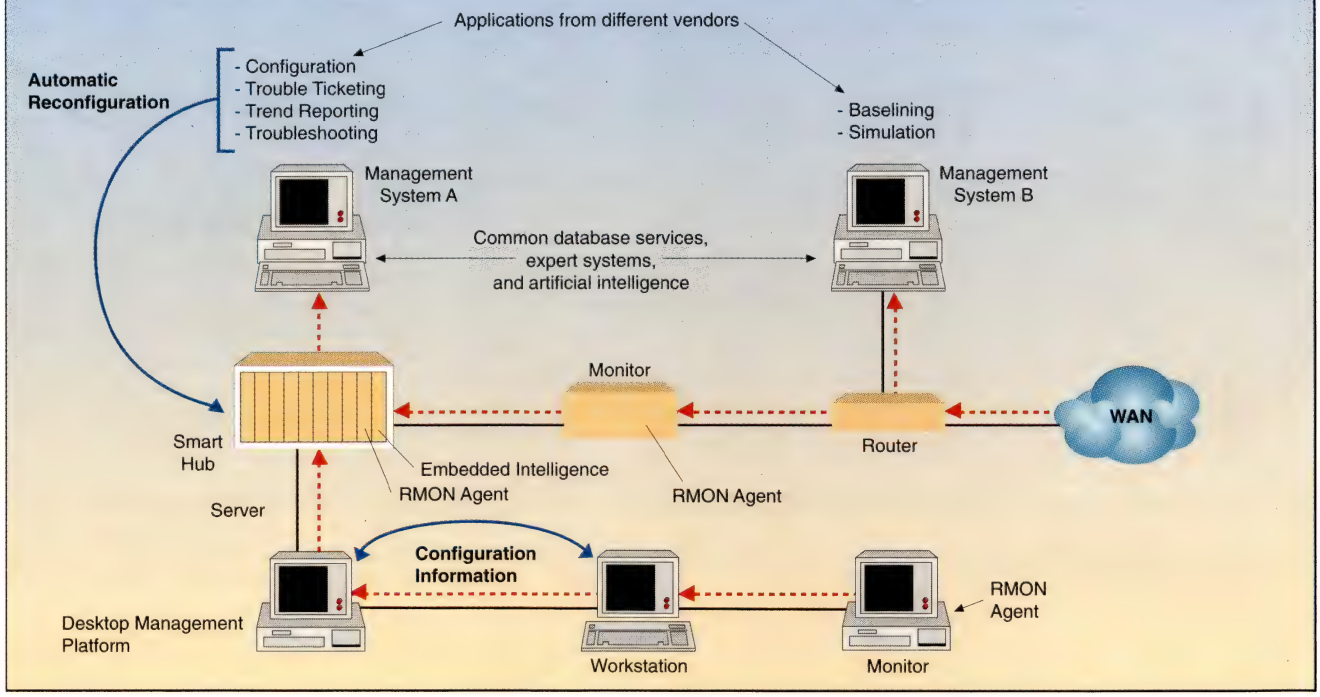

automatically or issue trouble tickets without returning to the main menu or opening a window. And the data created by such applications must be automatically added to LAN configuration files.

Vendors claim that their LAN management systems can do all of the foregoing, but many third-party developers disagree. They say that development tools for management systems are limited, supplying access only to a system's SNMP capabilities. This makes it possible to monitor devices on TCP/IP networks, display information on a map, and change an icon's colour if trouble occurs, but getting data from networks that aren't managed by SNMP is difficult and, at times, impossible. And even if it's possible to obtain all the data, getting management applications to share it remains a challenge.

Even supporting a common relational database (offered as an option by all leading vendors of management systems) means translating data into a format the database can use. Often, the APIs needed are available only from the database vendors.

To simplify matters, platform vendors are rewriting their system software in object-oriented formats and streamlining their APIs. For example, Hewlett-Packard recently announced it was working on a shared relational database that could be used by all applications running on OpenView. Only one set of APIs will be needed, and vendors will be able to access any data handled by the system, whether it comes from TCP/IP, NetWare, SNA, or DECnet.

IBM also says it has recast and consolidated its NetView/6000 APIs to make them more efficient. The company says developers can write applications that address data in various formats, including flat file or relational, via one set of APIs.

DEC too has enhanced its Polycenter Framework (formerly known as DECmcc) to make it easier for end-users and developers to add or integrate applications. DEC now offers software that lets users incorporate their Unix, VMS, or NetView scripts in Polycenter Framework applications.

Meanwhile, third-party developers are also seeking solutions. The long-awaited Distributed Management Environment (DME) developed by the OSF should also help create interoperable network management applications for Unix systems running the OSF's Distributed Computing Environment (DCE). The DME object-oriented software and development tools will allow applications to share data easily. DME code is slated to ship by the end of 1993.

The Desktop Connection

These integration efforts also will help bring desktop management into the realm of proactive tools. Desktop management products monitor servers and workstations on networks. Systems from IBM, Intel, Microcom, Network Computing, Novell, and others track network activity, disk space, input/output, and other system-level activities.

Some vendors bundle all their tools into one package. Intel's LANdesk, for instance, provides LAN monitoring, statistical analysis, and workstation and server system management functions. Others sell their management components separately.

To make it possible for proprietary platforms to work together, a group of vendors has founded the Desktop Management Task Force and is developing the Desktop Management Interface (DMI). The membership includes DEC, HP, IBM, Intel, Microsoft, Novell, Sunconnect, and SynOptics.

DMI, slated for release soon, will enable desktop management applications to run in a common memory space on a PC and communicate with management systems that support SNMP, CMIP (common management information protocol), or other standard protocols. It includes a component interface and a common database of definitions for desktop system components (all hardware, software, and peripherals in a PC) called a management information file (MIF). The component interface translates data from the MIF into the format used by the management protocol. This means that PCs and other workstations will be able not only to house multiple desktop applications, but also to access standard management platforms like HP OpenView.

Some vendors will integrate desktop management, including DMI, into their specific operating systems. Microsoft says that DMI, SNMP, and NetView will all be accommodated in Windows NT. Microsoft also plans on delivering its own software administration and LAN inventory applica-

MORE IN LESS

CONTROLLED ACCESS UNIT

Meet the Olicom CAU—an intelligent hub with more features packed into a mere two and a half inches of rackspace than anything else in the field.

Auto-bypass and disabling of defective lobes, nodes and components, together with out-of-band management, provide unsurpassed network availability. And outstanding access control and management facilities guarantee top security. What's more, it's the only CAU in the world that supports both IBM LAN Network Manager/NetView and SNMP.

The Olicom CAU is part of a full range of Token-Ring CAU cabling products—including STP LAMs, UTP LAMs and CAU Fiber Converters. They're all 100% IBM compatible and interoperable, and fully support IEEE 802.5 Token-Ring standards.

Take a closer look at the CAU that's putting the competition out to pasture. Call Olicom today.

olicom

TOKEN-RING SOLUTIONS

Australian Distributor for Olicom: FORCE TECHNOLOGY PTY LTD, P.O. Box 1516 · Dee Why, NSW 2099 · Ph: (02) 971 1000 · Fax: (02) 981 1932

Putting Network Monitoring at EASE

A compelling technology for gauging network baseline performance is supplied with the Ethertwist hubs and bridges from Hewlett-Packard's Roseville Networks Division, based in Roseville, California. These devices are equipped with the vendor's Embedded Advanced Sampling Environment (EASE), a series of algorithms in firmware that automatically determine normal network behaviour by continuously sampling traffic and comparing normal and abnormal parameters.

HP plans to offer a trio of Unix applications for EASE-equipped hubs. The Openview Traffic Expert will analyse data supplies by EASE and recommends actions, such as adding a bridge or router to create a network segment. The OpenView History analyser will evaluate historical data gathered by EASE from multiple LAN segments and determine traffic trends based on utilisation to assist in capacity planning. The OpenView Resource Manager will monitor traffic patterns through the transport layer, gauge normal behaviour, and flag abnormal events.

The Traffic Expert will watch for congestion in a hub based on what its algorithms have determined is normal saturation. If congestion is found, it will send the network manager a simple, easy-to-understand message, such as 'Move Sebastian's PC from Marketing East to Marketing West.' And the Traffic Expert can also tell managers when and where to place segmentation devices to improve LAN performance.

The History Analyser will be able to track network usage over time to establish trends. It might, for example, reveal that a CAD department is generating 60% of a company's X-Windows traffic over a six-month period, or that an order-entry division generates 50% of its electronic mail.

The Resource Manager will let users home in on problems. For instance, in a graphical view of network statistics, a user can click on an axis that shows an abnormal increase in packet activity. This brings up another view of the exact building or site where the anomaly is located (such as a multicast storm on a particular segment). Clicking again will take the viewer to the exact workstation that is causing the problem, showing its address and indicating what other stations it's talking to.

Hewlett-Packard says that EASE delivers the continuous filtering found on a LAN monitor without requiring users to pay for one. HP claims that EASE goes beyond monitoring, not merely identifying problems on multiple segments but also isolating the nodes that cause them. EASE does this by reading the network addresses inside errored packets.

EASE also tracks activity on the connections between network devices, which is beyond many LAN monitors. (Notable exceptions include tools from Concord Communications, Frontier Software Development, and Metrix Network systems.)

HP has added EASE to its Ethertwist Hub and Ethertwist Hub 48, as well as to its ThinLAN and Fiberoptic Hub Plus line.

HP needs to sort out its marketing strategy regarding EASE and its other LAN tools. EASE is touted as a continuous monitor, as is the LANprobe II monitor. EASE applications advise network managers, just as the Network Advisor protocol analyser does. Meanwhile, rumours persist that EASE will be combined with the LANprobe II and the Network Advisor. HP also said that talks were under way with respect to plans to license EASE to vendors of monitors and hubs.

Mary Jander

tion, code-named Hermes, for NT. Hermes will run under NT but will monitor systems under LAN Manager, LAN Manager for Unix, DEC Pathworks, Novell NetWare, IBM LAN Server, and IBM LAN Manager for OS/2. Hermes will automatically create a DMI management information file that will enable it to manage workstations and servers running under NT.

Hermes also will incorporate an SQL Server database for storing configuration information along with data about applications, job specifications, and audit trails. According to Microsoft, any hub, bridge, router, printer, or fax board can add itself automatically to the database by presenting a binary version of a DMI MIF for Hermes to collect. This means that proactive network management applications can integrate their data with the Hermes SQL database. This would be done using APIs, such as dbLib, an API offered by Microsoft for SQL Server, WOSA ODBC, or other NT APIs.

Real Simulation

Sophisticated simulation tools are available from vendors like US-based firms CACI Products and Comdisco Systems. These run on PCs or workstations and use mathematical techniques to mimic the actual performance of a LAN.

Integrating such tools with LAN management products like monitors and protocol analysers is a priority for simulation vendors, since incorporating actual information into a LAN simulation will make it more accurate. CACI and Comdisco allow users to transfer into their simulation programs data captured by Sniffer LAN protocol analysers from Network General. But this requires manually adjusting the analyser samples to run in the simulations; in the future, this link will become increasingly automated.

Simulation vendors also want to build direct links between their products and management systems. CACI announced in February that it plans to integrate its Comnet III object-oriented LAN/WAN simulation package with IBM's NetView/6000. This will make it possible to use information gathered by the NetView/6000, including LAN performance, fault management, and configuration data, in Comnet III simulations. Other simulation vendors say they too are planning alliances with net management vendors.

Artificial intelligence and expert systems also are emerging as potential tools for streamlining LAN management and making it more proactive. Two competing protocol analysers, the Expert Sniffer from Network General and the Network Advisor from Hewlett-Packard, offer users step-by-step advice for solving LAN problems. If a Token Ring LAN reports a station error, for instance, the HP Network Advisor will list the steps a LAN manager should take to determine whether a LAN adaptor is running at the wrong speed. And the LAN Command Advanced desktop management system from US firm Dolphin Networks uses an expert system to analyse Novell NetWare servers and workstations.

The AI Angle

Artificial intelligence could also be used to anticipate problems that may lead to network crashes. When a bridge fails on a network, for example, it's likely to affect the segments and devices attached to it. AI software could automatically swap in a backup bridge and reroute traffic accordingly, preserving communications to the imperiled segments.

Artificial intelligence can also be used to speed up the process of sifting through multiple alarms, which can be generated when one failure triggers another and all trouble indicators are lit at once on the main display of a management system. US firm Applied Computing Devices recently introduced an application called Correlator that uses expert system and neural networking technology to isolate the root of a problem when a management system generates a flurry of related alarms.

Mary Jander is network management and new products Editor for US networking magazine Data Communications.

Where in the world do you turn for internetworking solutions?

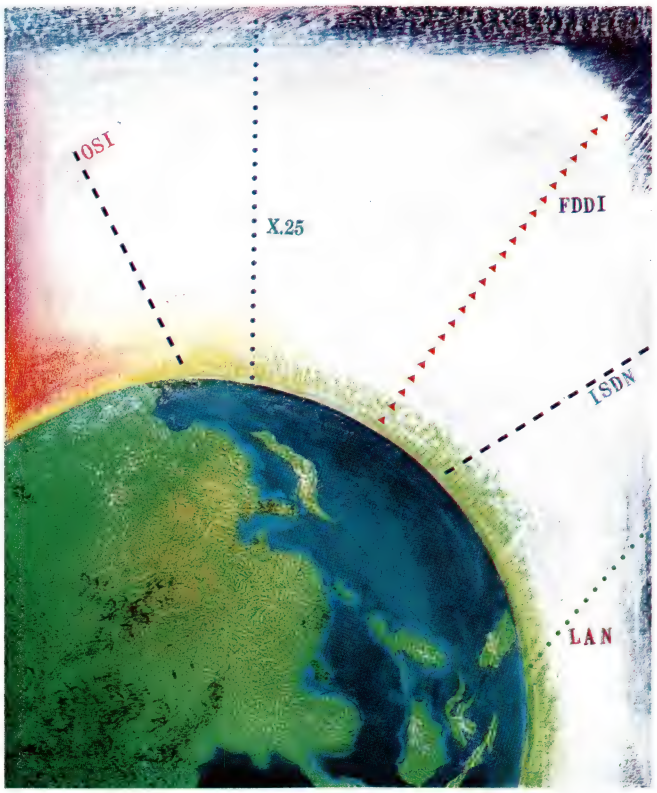

Today's multifaceted information networks require the world's most advanced diagnostic systems for protocol analysis and simulation. But where can you find the expertise necessary to fully implement and test WANs, ISDN, LANs and FDDI internetworking?

All paths lead to Tekelec.

More and more telcos, communications equipment manufacturers, and network managers worldwide are turning to Tekelec for leading-edge solutions that support corporate internetworking. Tekelec systems analyze, monitor and simulate networks. Reduce development costs and downtime. Enhance network integrity and quality. And ensure interoperability and conformance to emerging global standards.

A world of protocol experience.

From the R&D lab to the field, Tekelec's Chameleon line of protocol analyzers support the development and application of advanced data protocols like ISDN and OSI.

As the industry standard protocol analyzer and simulator for R&D and evaluation labs, the *Chameleon 32* supports all major protocols and interfaces, offers comprehensive network performance statistics, and features extensive simulation capabilities.

The *Chameleon 20*, a smaller version of the Chameleon 32, is ideal for both field service and network center applications.

Designed for field service, the *Chameleon 5* provides protocol analysis, trunk testing, and data line monitoring in one lightweight, rugged unit.

And our new *ChameLAN 100* gives product developers and network managers complete monitoring and simulation of high-speed FDDI networks.

Talk to Tekelec.

With the proliferation of WANs, ISDN, LANs and FDDI, the world is turning to Tekelec for advanced diagnostic systems. Call or write today for our new 12-page product overview brochure. And discover why the world is turning to Tekelec for internetworking solutions.

TEKELEC

Tekelec Australia Pty Limited
Suite 14, West Court, Bankstown City Capital Centre
41-45 Rickard Road, Bankstown NSW 2200 Phone: (02) 790 2200 Fax (02) 790 0848

Stay on line in the world of Office Technology.

■ Mobile phones ■ Business Communications ■ Network systems ■ Office equipment

THE AUSTRALIAN INTERNATIONAL COMMUNICATIONS AND OFFICE TECHNOLOGY EXHIBITION

ROYAL EXHIBITION BUILDING MELBOURNE 10-13 AUGUST 1993

TUES 12 NOON – 7.00 pm WED 10.00 am – 7.00 pm THURS *LATENIGHT* 10.00 am – 9.00 pm
FRI 10.00 am – 5.00 pm

CHILDREN UNDER 18 NOT ADMITTED. ADMISSION $14

Staged in conjunction with PC93 – Australia's leading computer exhibition

aes Australian Exhibition Services Pty Ltd 424 St Kilda Road Melbourne Vic 3004 Australia Tel + 61 3 867 4500 Free Call 008 335 010 Fax + 61 3 867 7981 Telex AA 39329

Multimedia Router

SynOptics has introduced a multimedia router for use in remote distributed offices. Company officials say the new router delivers high performance at a better price than backbone routers.

The new Model 3800 Multimedia Router can link multiple Ethernet, FDDI, Token Ring and wide area networks. Additionally, SynOptics announced the Model 3809 FDDI Personality Module which, when installed on Model 3800 Multi-MAC router baseboards, adds an FDDI network connection option.

The new product is based on a joint development between SynOptics and Cisco, and is compatible with Cisco's family of routers. The Model 3800 supports up to 16 simultaneous routing protocols, which SynOptics says provides efficient operation in all protocol environments. The protocols allow routers to exchange information to determine the optimal path between any two network segments. The company says the unit can also be configured as a full multiprotocol router, or can serve as a full function bridge where routing is not possible, or bridging is preferred.

Pricing for the Model 3800 Multimedia Router Base Unit with one Ethernet Personality Module is $12,145. The Model 3800 with one Token Ring Personality Module costs $12, 982. Personality modules can also be purchased separately.

SynOptics (03) 853 0799

Token Ring Software

Olicom has announced new software for its Token Ring bridges and local area network adaptors, which a company spokesperson says will provide a two-fold boost in performance over its existing products.

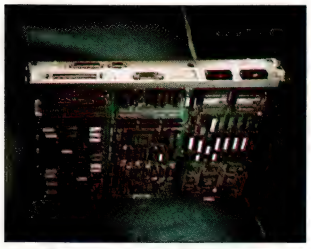

SynOptics' Model 3800 supports up to 16 simultaneous protocols

The new PowerMAC, when used with Olicom's standalone Wire-Speed 16/4 Token Ring Bridge, enables users to achieve throughput of 27K packets per second (Kpps) — a substantial increase over the existing speed of 13Kpps.

Company officials said the PowerMAC driver, used in conjunction with Olicom's EISA adaptor cards, lets PC users running NetWare boost throughput without putting increased demands on the CPU, achieving speeds of 19Kpps for 64-byte packets. Existing users of PowerMAC software can upgrade to the new software free of charge.

Force Technology
(02) 971 1000

Intel SatisFAXtion

Intel has released a new generation of PC facsimile products which are designed to allow PC and network users to bypass facsimile machines.

The new Intel SatisFAXtion Board includes an 80186 processor with 128K of RAM and a 2,400bps MNP5 modem. The board complies with the Communicating Applications Specification (CAS), and can produce 'near laser quality' output for Windows documents and high quality dot matrix printer quality output for DOS documents. It also provides genuine background fax processing, terminate-and stay-resident operation, automatic redial and a fax scheduler, company officials claim.

Priced at around $1,200, the Intel SatisFAXtion Board also features a port for a hand scanner to allow users to include images that do not originate on a PC in facsimile transmissions.

NET SatisFAXtion Version 2.0 is designed for use with Novell NetWare LANs and supports up to 1,000 users and up to eight Intel faxmodems per server. Company officials said the software can be automatically installed by users, has NetWare bindery import capabilities, an enhanced set of faxing features and support for Intel's SatisFAXtion board.

Users can be set up remotely across the network without the

HP's new NetServer LM Series is available in three models

HP Tilts at High-End Server Market

Hewlett-Packard has released two families of network servers that company officials say blend performance and expandability. The new HP NetServer LE Series is based on the Intel 80486 processor, while the high-end HP NetServer LM Series utilises Intel 80486 and Pentium processor technology. Both systems can be managed by HP's new HP NetServer Assistant server management tool.

The NetServer LE Series is designed for small to mid-sized workgroups and is available in three standard models based on 33MHz 486SX, 33MHz 486DX and 66MHz DX2 processors. Each model can be upgraded to a Pentium overdrive processor, say company officials. Supplied with a three year warranty, the product family can support up to 128MB of RAM, 256K write-back cache and integrated Fast SCSI-2, IDE and Super VGA controllers. Pricing begins from $3,651 (excluding tax) for a base model; $5,290 (including tax) with a 33MHz 486 processor and 240MB hard drive; and $8,085 (including tax) for a NetServer LE based on a 66MHz 486 DX2 processor and 535MB hard drive.

Featuring a new 'Power Cabinet' design, the high-end HP NetServer LM Series is available in three standard models based on 33MHz 486SX, 66MHz 486DX2 and 60MHz Pentium processors. Each model takes advantage of an integrated 32-bit fast SCSI-2 controller and a large secondary memory cache, company officials say. In addition, the servers feature flash EPROM to simplify system updates and compatibility with redundant arrays of inexpensive disks (RAID) technology.

To boost performance to higher level, the NetServer LM Series can accept a dual Pentium processor board upgrade for symmetric multiprocessing, while upgrading of 486-based systems to a Pentium processor is facilitated by a special zero insertion force (ZIF) socket on the system board.

Each NetServer LM is supplied with 16MB of RAM, expandable to 384MB; a 3.5 inch 1.44MB floppy drive and front accessible, half height storage shelves capable of handling up to eight mass-storage devices. Bundled with HP NetServer Assistant, indicative (inclusive of tax) pricing begins at $9,685 for a 33MHz 486-based system with a 535MB hard drive; and ranges to $14,170 for a 60MHz Pentium based system with a 535MB hard drive.

Built on HP's OpenView network management platform, HP server Assistant helps identify and prevent server problems, company officials said. Designed in its first release to support Novell NetWare, the software makes it easy to integrate third party network utilities and provides a number of tools to resolve network problems and allow remote network management.

Hewlett-Packard 13 13 47

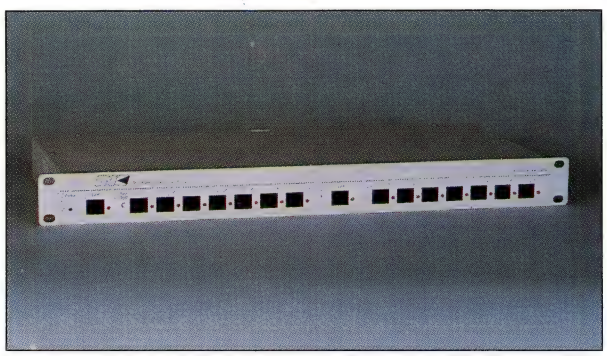

INC's 5591/5592 connect PCs to IBM AS/400 hosts via twinax

INC Releases Active Star Range

INC Manufacturing, the wholly Australian-owned and operated manufacturer of local area network connectivity products for IBM AS/400, Token Ring and Ethernet networks, has announced the releases of its Active Star 5250 product range.

Available in 7- and 14-port models, the new INC 5591/5592 and INC 8594 Active Stars are designed to allow fault tolerant PC connection through twinax cable to IBM AS/400 hosts. Company officials claim the devices accurately reconstruct signals to their original state using retiming, reclocking and reshaping techniques.

Featuring front panel two-colour LEDs to allow error status monitoring, the products include a self test switch and support link and drop distances of up to 1,500 metres (twinax and shielded twisted pair cabling), and 1,000 metres (unshielded twisted pair cabling). The INC 8594 uses IBM Type 1 shielded twisted pair cabling, while the INC 5591/5592 is available for use with RJ45 unshielded twisted pair cabling.

Available now from INC distributors Anixter Australia, MPA International, Hardie Networks and LAN Systems, prices (excluding tax) range from $1,550 for the 7-port INC 5591; $1,800 for the INC 8594; and $2,400 for the INC 5592.

INC Manufacturing (02) 525 8411

administrator needing to set up software at each workstation, while NetWare bindery import eliminates the chore of creating a fax user database name by name. Existing phone book databases (in ASCII or CSV format) can also be imported into the NET SatisFAXtion phone book.

Also supported is semi-automatic inbound routing (SIR), which forwards faxes to the fax administrator for routing after viewing only the cover page. DOS or Windows alerts can be used to signal incoming faxes to the fax administrator.

NET SatisFAXtion and the Intel SatisFAXtion Board also incorporate Intel's WYPIWYF (what you print is what you fax) interface, which allows users to send faxes from within a DOS or Windows application by using a print command. Support for HP's PCL5 scalable font technology helps provide a rich set of fonts for outgoing faxes, they said.

NET SatisFAXtion Version 2.0 software is priced at around $1,500 for 20 users and $2,500 for 1,000 users.

Com Tech (02) 317 3088

Remote Bridge/ Concentrator

ADE Network Technology has released the ILC/200 Remote Bridge/UTP Concentrator.

Company officials say the product combines a full E1 remote bridge and a 25-port concentrator into a single unit, providing network users with a low cost 10Base-T network at a remote office.

The unit is designed to be used in multiprotocol, multivendor environments, and provides SNMP agent support, and can be monitored and managed by SNMP-based network management stations from other system vendors. Users can get statistics on readable frames, short events, collisions, late collisions, data rate mismatch and auto partitions, and individual ports can be enabled and disabled remotely, said officials.

The ILC/200 also features SNMP bridge agent support, spanning tree protocol, custom filtering and collection and reporting of bridge system status and statistics.

ADE Network Technology (03) 543 2677

New Repartee

US-based Active Voice Corporation has announced a new version of its Repartee Voice Processing System. New features include a menu interface, speed keys, special delivery options, and subscriber self-enrollment.

The new Repartee comes with voice mail, automated attendant and audiotext, and is the only system that allows users to customise their menu options, according to a company spokesperson.

Repartee is available with 2 to 36 ports and up to 90 hours of storage. Standard features include: automatic setup on all major telephone systems; message receipt verification; day and time stamp; security codes; group broadcasting; user-changeable voice prompts; numeric access; voice detect; and personal greetings.

Active Voice (03) 894 1699

SNMP Site Manager

Wellfleet Communications recently announced the release of its Site Manager SNMP-based node management application for DOS-based 386 and 486 PCs running Windows.

The product was previously available only for Sun SPARCstations running Sun Open Windows or Motif.

The Site Manager is designed to enable Wellfleet bridge/routers to be configured, monitored and controlled over the network from local and/or remote workstations, said company officials.

It provides a comprehensive set of network node management capabilities that include: centralised configuration management; real-time performance and error analysis; and real-time event and fault monitoring to simplify problem identification and isolation.

Wellfleet (02) 959 1290

Multi-link Datacom Analyser

Australian companies Testcom Data and Mobvan have developed the System 9000 MLDA, which allow simultaneous exercising of up to 12 data communications links running at speeds from 300bps to 2Mbps, using a variety of interfaces.

Company officials said the product had been developed in response to a request from Optus Communications for a high-density BERT system that could test several links at the same time and be controlled from a central console.

The new system is housed in a standard 19-inch rack-mount unit and contains from 2 to 12 cards, which are fully featured, high speed BER testers equip-

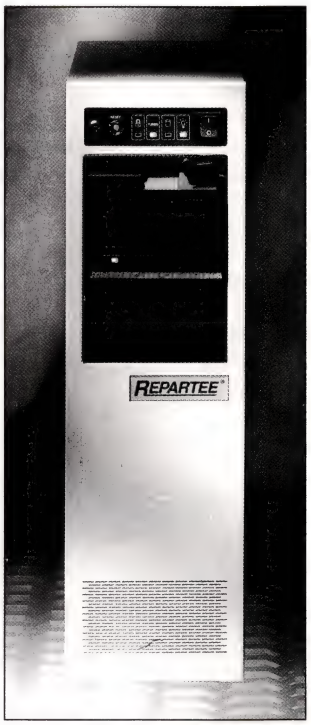

Repartee lets users customise their menu options

Stallion's ETS can connect to previously incompatible environments

ped with all common interfaces (for example: V.24; V.35; X.21 (RS449); G.703/64Kbps co-directional; G.803/2.048Mbps).

Both internal and/or external clocking allows the new System 9000 MLDA to operate at data rates of 300-19.2Kbps async and 600-2.048Mbps sync, said a company spokesperson.

Users can select a variety of test patterns, and test results are presented in BER and G.821 format. Alarm messages are displayed, and signals sound when errors exceed the pre-set levels.
Testcom-Data (02) 630 7528

Ethernet Servers

Australian Unix developer Stallion Technologies has announced two new products — a terminal server, called the ETS, and a communications server, called the ECS. The products form part of the company's Easy Server family, and officials claim they provide secure and open access connectivity for serial devices to multiple host systems on TCP/IP and DEC LAT networks.

They said one of the main benefits of the products is their ability to connect previously incompatible environments. They support concurrent TCP/IP and LAT, SNMP, RIP and ICMP, as well as standard features such as Telnet, Rlogin and ping.

For installations with more complex requirements, ECS has additional features such as TCP to LAT gateway services, higher level network security, and the ability to centrally manage multiple servers.

The ability to provide connections to remote devices and LANs via low-cost serial lines and asynchronous modems is provided via the Serial Line Interface Protocol. A compressed version (CSLIP) and Point-to-Point Protocol (PPP) are also supported.

Company officials said both products support rotaries, allowing printers and modems to be pooled, which provides a load balancing scheme which lets users access more lightly loaded system resources. Also supported is reverse Telnet, allowing Unix programs the use bi-directional communications to run unmodified, transparently connecting users to printers and modems across the network.

The new servers also support Univel's UnixWare, SCO Unix and Xenix, Unix SVR4 and Interactive (SunSoft) Unix. Both the ETS and ECS are available in 8 and 16-port server modules.
Stallion Technologies
(07) 870 4999

Network Management Extender

Proteon's new Network Management Extender provides centralised management for up to 16 LANs from a single device, according to officials from the company.

The Extender allows up to eight Token Ring and eight Ethernet local area networks to be managed via a single device, thereby eliminating the need to buy multiple network management modules, according to officials from the company.

The new product supports all SNMP-compliant management packages, and condenses all network information for transfer to any SNMP-based manager.

Other features of the product include the automatic mapping of MAC addresses to physical ports, the automatic recovery of persistent beaconing, and SNMP event logging.

The company said the Network Management Extender is especially suited to a head office environment which utilses a collapsed backbone architecture. Pricing for the Network Management Extender is $15,000.
Proteon (02) 955 8555

Mobile Fax Machine

NEC has announced what is says is the world's smallest mobile fax machine. The new i300 Image Terminal is the size of an A4 page, and can connect to any of NEC's current range of mobile phones. It weighs in at 2.1 kilograms, with a 0.8Kg battery, and comes in a leather attache case which has room for a mobile phone.

Company officials said the i300 is a full featured fax which is designed to operate as a free-standing, battery-powered mobile device when connected to a mobile phone, but which can also be connected to the public switched telephone network via an optional acoustic coupler attached to a standard telephone.

The unit features an error-correction mode which will retransmit corrupted data if noise or data corruption is detected. It also has an out-of-paper memory, and a NiCad battery that provides 18 hours standby time, or power to transmit approximately 50 pages.

NEC officials said the i300 retails for under $3,000, and will be used by workers in a wide range of occupations who need to transmit and receive documents in the field.
NEC 008 036 136

New External Hubs

Artisoft has introduced two new concentrators for use in 10Base-T installations. The models are the T-Runner 800/TC, which includes eight ports and is priced at $675, and the T-Runner 1200/TC, which has twelve ports and sells for $1,125.

Both are IEEE-compliant, and are compatible with a wide variety of networking environments, say officials. They contain a BNC connector that allows them to be linked to one another via thin coax cable, and also allows a mix of cable types to be used. The hubs also have an auto-partitioning feature that automatically disables ports that do not have link integrity.
Artisoft (02) 880 2688

Apple Extends Cross-Platform Capabilities

Apple has announced that it will extend the cross-platform capabilities of its client/server offer-

NEC's i300 Image Terminal is the world's smallest fax machine

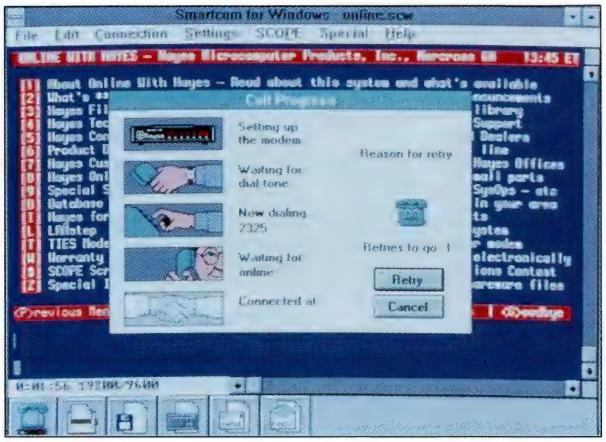

Hayes' Smartcom supports all major file transfer protocols

Sophisticated Windows Communications

Smartcom for Windows, a full featured asynchronous communications program for Microsoft Windows, has been recently released by US-based Hayes Microcomputer Products.

Priced at $245, the software combines an easy to use graphical user interface without compromising power and performance, company officials say. Included among its key features are: Smart Buttons, user-definable macro keys which can be customised with both icons and text; Smartcom for Windows Scripting Language (SCOPE), which allows users automate tasks and create interface-driven scripts to simplify complex actions; context-sensitive help; DDE support; and Smartcom for Windows Communications Editor, a built-in word processor with ANSI.SYS functionality.

With a range of built-in terminal emulations, the product also supports all leading file transfer protocols, including: Xmodem (standard, CRC, 1K and Ack-Ahead), Ymodem, Ymodem-G, Zmodem, Kermit and Compuserve B Plus File Transfer. In addition, the software supports an 'autotype' protocol which allows users to automatically 'type' a disk file, the clipboard or contents of the Communications Editor in either ASCII or ANSI.SYS format to a remote computer.

For network users, the product supports both the NASI/NCSI interface and the Int 14 interface which combine compatibility with communications servers such as Novell NACS and Ungermann-Bass NETCI Servers. TCP/IP Telnet communications are supported, as are individual standalone shared modems or multiple modems on asynchronous communications servers.

The product also boasts support for multiple languages, including: English, French, German, Italian and Spanish. More than one language can be installed, and users can quickly switch languages, company officials say.

MPA International (03) 724 4444

ings by making it possible for many Macintosh software services and application software to run on Unix systems.

Apple said it will work closely with leading companies and standards organisations to provide Apple technologies on their Unix servers and clients. Apple will provide an open, published set of APIs which will allow Mac and Unix developers to enhance the functionality and ease of use of their products for client/server computing, according to an Apple spokesperson.

Apple (02) 452 8000

New IBM LaserPrinters

Lexmark has introduced the new IBM LaserPrinter 4039 family of intelligent, high-resolution printers that offer advanced networking capabilities.

Pricing for the five models ranges from $2,995 to $5,950, and print speeds range from 10 to 16 pages per minute. All models feature a high performance AMD RISC processor, built-in PostScript and PCL5 compatibility, and a SmartSwitch function that automatically selects the right language for the print job, say officials.

The printers provide 600 dpi resolution, and feature an array of intelligent capabilities which enable them to work with almost every kind of commercial software application, according to Lexmark officials.

The 4039 family also features flash memory, which stores frequently used forms, custom fonts and logos. All models can support direct network connectivity and bi-directional communications with the optional MarkNet family of Internal Network Adaptor cards, said Lexmark officials.

A 4-line by 20-character liquid crystal display allows users to interact with the printer in simple English, or any of seven other languages.

Lexmark (02) 930 3500

New SCO Operating Systems

The Santa Cruz Operation has announced the introduction of a new family of server and workstation operating systems especially designed for the client-server market.

The two main product lines are: SCO Open Server Release 3.0, and SCO Desktop Release 3.0. The company says the new products offer several advantages, including:

- All new products are based around SCO UNIX System V/386, Release 3.2, version 4.2, which offers better performance, configurable security, increased system capacity and simplified system management;
- Better packaging to suit user requirements;
- An X Window X11 R5, and Motif 1.20 environment;
- New versions of NFS, TCP/IP, OSI and SCO MPX;
- Windows 3.1

The new family also features a new configuration called SCO Open Desktop Lite, designed for lower end, mobile systems.

SCO (02) 911 7711

Fault Tolerant ERS

Cubix has released a fault tolerant version of its Enhanced Resource Subsystem, called the ERS/FT.

The system includes redundant hot-swappable power supplies, and the ability to power-down an individual communications processor without affecting the rest of the unit.

The new ERS/FT uses a segmented EISA or ISA passive backplane in conjunction with Cubix's plug-in CPU cards, and creates a network platform that supports up to eight fully functional 386 or 486 communications or applications servers, according to officials from the company. They added that users can segment the backplane in a variety of ways to create custom platforms, and can hold up to four systems in a rack-mount enclosure.

Wholesale Technology
(02) 368 4800

2Mbps WAN Analyser

Hewlett-Packard has announced the HP 4959B WAN protocol analyser, which allows technicians to troubleshoot WANs operating at full utilisation at 2-Mbps transmission rates.

A HP spokesperson said this is an advantage, because it is in this operating mode that most WAN problems occur.

The unit's data capture hardware ensures data integrity at maximum network utilisation,

The 4039 printer family offers high-level networking capabilities

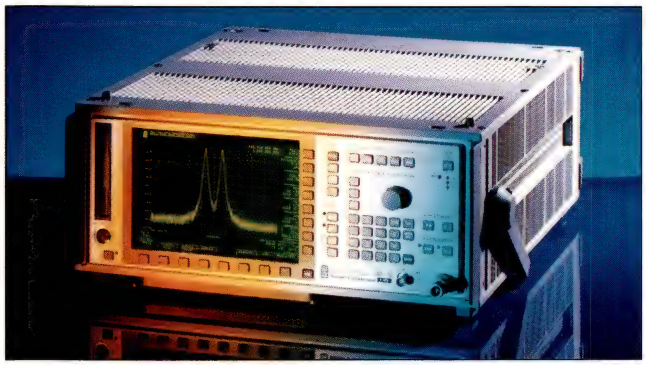

Test data captured by Wandel & Goltermann's SNA-23 microwave spectrum analyser can be downloaded to a PC for further analysis

which enables problems to be pinpointed accurately.

The new HP 4959B allows users to select configurations that automatically set up data filters, triggers and counters for LAN interconnection tests. The unit has a complete range of functionality, from statistical analysis, to data decodes, simulation and bit error-rate testing, company officials.

It supports a variety of high-speed WAN technologies including X.25, frame relay, SMDS, ISDN (both basic and primary rates), IBM SNA and most synchronous and asynchronous protocols. Internet data analysis is supported for the most popularly used local area network protocols, including IP(1) and SNAP(2).

The company said the HP 4959B supports RS-232; RS-449; V.35, T1 (1.544Mbps) or E1 (G.704-2.048Mbps) hardware interfaces. It comes with a standard 4MB data capture buffer and an 86MB hard disk drive, and captured data and screens can be sent to an external printer over an RS-232C connection.

Base price for the unit is $32,000.

Hewlett-Packard 13 13 47

Portable Microwave Analyser

Wandel and Goltermann have announced a new microwave spectrum analyser designed especially for field testing of radio systems.

The new SNA-23 has a frequency range from 9kHz to 26.5GHz, and offers high sensitivity, wide dynamic range (approximately 90dB, even at the higher frequencies), and excellent frequency accuracy, according to officials from the company.

The W&G SNA-23's user interface and application-oriented test functions make it well-suited to testing radio link and satellite communications systems, the officials said.

Its main applications are field testing of radio systems, rapid detection of low amplitude interference, and the taking of measurements where intermodulation cannot be tolerated.

The unit has a large electroluminescent display. Printer and plotter interfaces come standard with the unit, allowing users to easily document results. Stored results can also be downloaded to a PC for further processing.

Wandel & Goltermann
(03) 690 6700

Network Management Software

Unixpac has announced Release 1.6 of the Tivoli Management Environment from Tivoli Systems.

The software is designed to help organisations with computing resources dispersed across networks of servers and workstations to reduce substantially the cost and complexity of managing those systems, a Unixpac spokesperson said.

The Tivoli Management Environment provides system administrators with a set of software products that automate and simplify the repetitive processes required by enterprise-wide networks, such as installing software, adding new users and workstations, maintaining security, and monitoring key system resources, the company spokesperson said.

Scitec's IDM 300 can be tailored to a company's specific needs

The product has been designed for use on workstations from Hewlett-Packard and Sun Microsystems running Unix. Release 1.6 incorporates new technology, called configuration and change services, that automates the process of setting up, maintaining and changing enterprise-wide networks — work that typically occupies around 50% of a system managers' time, according to the spokesperson from Unixpac.

The product is available immediately for Sun Microsystems workstations and servers, and will be available for Hewlett-Packard systems in the third quarter of 1993.

Pricing for Release 1.6 starts at $39,000 for a 20-node network, which includes the Tivoli Management Framework, Tivoli/Works, Tivoli/Courier, Tivoli/FSM, Tivoli/Sentry and Tivoli/AEF.

Unixpac (02) 953 8366

Digital Multiplexer

Scitec has announced a new low-cost intelligent digital multiplexer for companies with 2-Mbps links, or those wanting to use the new fractional E1 services that will be offered by the carriers.

The IDM 300 is a full digital cross connect switch, with an 8-slot modular architecture, which is incrementally upgradable and can be tailored to suit a company's needs.

Company officials said it is specifically designed for LAN interconnect, PABX links, integrated voice and data, video conferencing, and CAD/CAM. It supports G.703/704, RS422 and V.35 interfaces, and any port or timeslot can be cross connected to any other timeslot or port. The IDM 300's cards are hot-swappable, and redundant power supplies can be used to improve the system's reliability, said the officials.

Scitec (02) 428 9555

New Canary Products

JNA Network Services has announced that it will distribute networking products from Canary Communications, including Ethernet transceivers, mini-hubs and media converters.

The company says the new product range will enable it to offer turnkey Ethernet network solutions. JNA will make available Canary's X-Series Ethernet transceivers, which allow a number of PC devices to connect to a single X-Series transceiver; H-Series mini-hubs, which allow multiple Ethernet devices to be interconnected without the need for a coaxial or twisted pair cable; low-cost UTP and multimedia hubs; a full range of Ethernet and multimedia repeaters; and a range of media converters which allow users to convert between thick and thin coax Ethernet cable and 10Base-T unshielded twisted pair.

JNA Network Services (02) 417 6177

XGATE 4.0 provides worldwide access to IBM midrange systems

X.25 Gateway Upgrade

Perle Systems has announced the availability of version 4.0 of its XGATE System X.25 gateway to IBM AS/400 or S/36-38 systems.

Officials from the company said the system provides access to IBM midrange systems located anywhere in the world, with better reliability, accessibility and performance than traditional remote midrange access techniques.

The system utilises the global X.25 data network, and features: support for various X.25 window and packet sizes for international use; automatic selection of packet sizes, allowing users to choose the packet size according to the network; and automatic setup of X.3 parameters for easier installation.

Perle's XGATE System also comes with improved PerleTalk software for PCs or notebooks, with improved features such as network log-on scripts for user-designed send/receive scripts, a phone directory allowing users to define a set of phone numbers, and a dialling abort facility.

The software can also terminate 'out of memory' directly from DOS, allowing unattended operation, and it incorporates a 'time out' facility so the system is not tied up when it fails to connect.

Other new features include full 5394 controller support, improved printing options, support for a faster 19,200 baud data speed, and the ability to load the Adaptor Handler into high memory, and better hardware.

Perle Systems (02) 416 0650

IBM LAN Adaptor

IBM has released the new LANStramer MC 32, which it claims is the world's fastest Token Ring adaptor. The product is a Microchannel adaptor which uses the full spectrum of the Token Ring bandwidth, from 4Mbps up to 16Mbps. It can process over 30,000 frames per second, and IBM says it will extend the life of existing hardware, as well as providing the speed for multimedia applications.

The LANStramer MC 32 features: new high-performance Token Ring technology with 32-bit busmaster capability, and which operates with full Token Ring media support for large and small frames; priority queuing, allowing multiple separate streams of information to flow through the adaptor; multiple group addressing, with up to 256 multiple group addresses available; support for both UTP and STP; server and client device drivers for NetWare, IBM LAN Server, DOS, Windows and OS/2; and Remote Program Load as standard.

The IBM LANStramer MC 32 comes with a five year warranty and is priced at 1,650, excluding tax.

IBM (02) 634 9111

The PF-47 is a DS3 analyser in a battery powered handheld unit

Bull SDM for NetWare

Bull Australia has announced the availability of its Software Distribution Manager (SDM) Version 3.0 for Novell NetWare LANs. The new version is an expansion of the existing SDM product for Microsoft LAN Manager servers, and allows for the easy electronic distribution and installation of software from a central point in an organisation.

The new product comprises a NetWare Loadable Module (NLM) on the server, and enables any package or data to be distributed and installed, including server upgrades, operating system upgrades, word processing and spreadsheet packages, all communications software or custom business applications, according to company officials.

They added that networks with OS/2 and NetWare servers can be serviced by a common master distribution server, and users are able to select from protocols such as SPX, Named Pipes, TCP/IP and X.25.

The SDM administrator's graphical user interface comes in a separate module that can be executed on any participating SDM node. The OS/2 Presentation Manager version is available immediately, and the Windows and Motif versions will be released soon, said officials.

Bull (02) 986 3622

Handheld DS3 Analyser

Wandel & Goltermann's PF-47 Handheld DS3 Analyser provides a traditional DS3 Bit Error Rate Test Set in a battery-powered handheld case. Officials from the company said the unit is able to perform both in-service DS3 monitoring and out-of-service DS3 testing, and because of its portability, is suitable for the installation and maintenance of DS3 transmission systems at all customer sites.

The PF-47's AutoSetup feature self-configures the unit's mode, frame and pattern to the received DS3 signal. The test results can be downloaded to a printer or personal computer, or viewed on the backlit LCD. Officials said over 70 simultaneous measurements and analyses, complete PRBS/Programmable/-Digital Word selection, and comprehensive error insertion are also provided.

Wandel & Goltermann
(03) 690 6700

New INC Offerings

Australian manufacturer INC recently announced a wide range of new products. They included the 4561/2 8-port multiplexer, which allows for the connection of up to eight IBM S/3X or AS/400 ports over a single fibre optic or copper trunk line, and the 4049B 8-port Device Multiplexer, which allows the connection of up to eight IBM3270 workstations to an IBM 3174 controller over coax or twisted pair cable.

The company has also announced five new Token Ring products: the TR288A Active MAU, which comes with power source, support for 16 devices and sophisticated noise filtering for large area coverage; the 5014 and 5262 4/16Mbps Fibre Optic Repeaters, which extend transmission distance to around 5,000 metres between adjacent MAUs; the new 5334 Repeater, which extends transmission distance on copper cable to 360m

at 16Mbps, and requires only a pair to complete a link; the 5013 2-port Token Ring Lobe Expander, which doubles the capacity of a single MAU lobe to support two workstations instead of one; and the new 5387 Transceiver, which can convert Token Ring signals from copper to fibre optic.
INC Manufacturing
(02) 525 8411

All-Purpose Repeater

ADE Network Technology has released the Netcor UCR507 modular multimedia repeater. The Netcor UCR507 can be configured with anywhere from two to seven ports, which themselves can be any combination of 10Base-T, coax, AUI or fibre optic. Officials said this makes the new repeater a flexible and low-cost solution. Pricing starts at $1,235, with a range of RAU add-on modules available.
ADE Network Technology
(03) 543 2677

Digital Voice Recorder

Adacel has introduced its DVR-1000 Digital Voice Recorder, which the it says is an alternative to existing analogue multi-track tape recording systems.

The system uses Digital Signal Processing (DSP) techniques, combined with disk systems and high capacity digital recording tape. Officials said the DVR-1000 is capable of recording, logging and replaying up to eight voice channels simultaneously, and may be controlled locally via a graphical user interface, or remotely, via another computer.

Officials said the unit uses sophisticated compression algorithms to allow maximum utilisation of storage media, with an instant recall capability of up to 10 voice hours old, and with a logging capability of 150 voice hours per tape. Other features include: record to disk and to optional DAT; replay from either disk or tape; simultaneous replay while recording; replay in both normal and compressed modes; side-tone to selected outputs; quick search capability; built-in test and diagnostic functions; and password protection.
Adacel (03) 596 2991

The Netcor UCR507 supports 10Base-T, coax, AUI or fibre cables

New Bridge/Routing Technology

Cabletron Systems has unveiled new technology which the company says will allow it to provide integrated bridge and routing capabilities throughout its range of standalone and modular hubs.

The technology will be delivered through a family of products known as Bridge Router Interface Modules (BRIMs), that will reside in new products that the company will announce in the next few months, officials said.The BRIM family will consist of bridge/router interfaces designed to provide connectivity to all standard network technologies, including Ethernet, Token Ring, FDDI, WAN and ATM, and will enable Cabletron's hubs to provide in-the-hub local or remote routing.

Cabletron officials said the new BRIM technology will sig-

The Osgate 700 can integrate disparate devices into a network

nificantly reduce the cost of deploying multi-segment hubs and wide area connectivity for remote offices or workgroups.
Cabletron (02) 878 5777

Network Management Gateway

Oscom has announced the Osgate 700 Network Management Gateway, which has the ability to integrate disparate devices into the network management environment.

A company spokesman said the new Osgate 700 family of generic and custom programmable models allows integration and management of a range of devices, including uninterruptable power supplies, process controllers, PABXs, alarm systems, security systems, air conditioning systems, and environmental control systems.

Device interfaces include RS -232, Parallel I/O, Analogue to Digital, and Digital to Analogue. Incoming I/O signals are connected via an industry standard I/O rack to allow standard Input and Output modules to be used, said the spokesperson.

The Osgate 700 can integrate remotely and locally attached devices via LAN, WAN or link-attached Network Management stations, and staff can monitor all critical devices in a centralised or distributed manner. Interfaces include IEEE 802.3, X.25 and RS-232, and the hardware and software components are designed for unattended operation and automatic restart, which ensures reliable operation, even in harsh environments, according to company officials.
Oscom International
(09) 481 3444

Business Key System

NEC has released what it calls a 'future-proof' digital business telephone system.

The new NDK9000 enables medium sized companies to customise and expand services as they grow, increasing capacity with hardware and software upgrades. Up to 32 lines in one office can be accessed through a combination of eight, 16 or 24 keys, or, alternatively, all lines can be accessed by a single key, obviating the problems that can arise when a busy office system has too many lines and not enough keys.

The new NDK9000 comes in modules providing up to 64 ports, which can be configured according to the company's need for a certain number of lines to the exchange, and number of required extensions.

Battery backup maintains power to the system, even in the event of a total blackout, and a text-messaging feature provides up to ten customised messages which will be displayed on any internal phone which calls an extension. The NEC NDK9000 can also be split to serve multiple tenants within a small office complex, so the cost of the system can be shared between businesses.
NEC 008 036 136

Gravel Access Roads, Multimedia Highways, and Convergence

Recently, a large American telecoms carrier has been airing advertisements on US television featuring a flurry of 'coming soon' applications. The ads show hand-held video terminals being used by people lounging on a beach, and so forth. Likewise, President Bill Clinton's multimedia guru, Mitch Kapor of the Electronic Frontier Foundation (EFF), has held forth at great length about the electronic highways that must link America if it is to have a bright economic future. This was a campaign theme of sorts in last year's US presidential election, and the Clinton-Gore team wisely courted the technologists and computer/multimedia industry leaders.

What is perhaps most interesting about the grandiloquent speeches promising this great future is that much of what is being promised already exists. As in most developed countries, the US is close to an optical fibre national grid. The problem rests in the 'last mile' or local loop. 'Fibre to the kerb' just won't quite do. And residential use patterns do not, or not yet, warrant the sort of expenditure required to put fibre from the kerb to the home.

This chicken-and-egg argument is not new. Indeed, ISDN evolved in some part as a solution to this problem. Whether one sees ISDN as a transitional stage to full fibre or as *the* solution, it would certainly enable the multimedia world to be brought to users far more cheaply. In fact, the Pacific Telecommunications Council's proposed Special Working Group on an Asia/Pacific ISDN/Datacomms Users Forum has generated a large amount of interest based upon innovative ISDN uses, like 'ISDN-by-satellite' for Pacific Island countries.

The success of ISDN may fuel the subsequent connection of fibre to the home. Certainly, this seems to be the ultimate goal of most telecom carriers throughout the OECD countries. Even in developing countries, laying new copper is not really seen as much of a long-term solution. If stop-gaps are needed, then go wireless. If one is going to put in however many lines, you might as well go all the way and put in fibre. This is the lesson learned throughout Asia. Also, the early reliance on cellular to relieve a creaky network leads to a comfort and familiarity with wireless telephony that then complements nicely an advanced multimedia wired network. In this way, Asia leads in the 'psychology' of seamless telecom service.

In most countries, even within Asia, regulatory frameworks lag far, far behind the technology-neutral full-service telecoms environment that is already emerging. Convergence is a reality, and the blurring of the lines between telecoms, cable, broadcast, and content provision is well under way. It was no accident that the President of Apple Computer was the keynote speaker at the 1993 US National Association of Broadcasters (NAB) convention in Las Vegas. The NAB is *the* meeting of the US broadcasting industry and attracted over 64,000 (yes, 64,000) delegates from virtually every broadcast organisation in the world. Every year, more and more telecoms or multimedia organisations are present. Likewise, here at PTC we've seen a large increase in the participation and membership of broadcasting organisations. Among our keynote speakers at PTC'93, the 15th Annual Pacific Telecommunications conference, was the President of the North American National Broadcasters' Association, Michael McEwen of the Canadian Broadcasting Corporation.

The recent US West-Time Warner deal shows clearly how this cross-pollination is now entering high-stakes levels. The psychological gulf that could be perceived between telecoms people and broadcasting people is lessening steadily. At a recent multimedia conference in Canada, I was impressed by the easy fluency with which multimedia professionals — those on the creative side — bridge content and carriage. Of course, gaps still exist. Some of the larger, more conservative telecom carriers were noted by their absence. Too many broadcasters are so enveloped by programming, cultural, and/or marketing concerns that they forget how rapidly their technological environment is changing. Likewise, too many telecoms professionals are too focused on technology for technology's sake and are too dismissive of the creative side of the new convergent environment.

Ultimately, convergence will never lead to the merger of Cannes with the IEEE, but the two have much more to do with one another than either realises, and this interdependence will continue to grow.

Convergence will be the manifestation of this growing interdependence. To quote from a recent article in the US weekly *The New Republic*: "The phrase '500 channels,' which often crops up in discussions of the (electronic) super highway, is best banished from future discourse. The preferred model is Internet, the global computer meta-network which encompasses various institutional archives and millions of personal computers. On Internet people can read and post messages on bulletin boards; send e-mail to a friend; 'broadcast' e-mail to lots of friends; search databases and download articles and software; play games by long distance; join in discussions — whether in real time or on bulletin boards — that become magazines in progress read by lots of people who don't contribute; and so on. Friendships blossom, clubs form, debates rage . . . It is with evident satisfaction that Mitch Kapor calls it 'one of the world's largest functioning anarchies'." [Wright, TNR 208:21, 24 May 1993].

"Ultimately, convergence will never lead to the merger of Cannes with the IEEE, but the two have much more to do with one another than either realises, and this interdependence will continue to grow."

As noted above, PTC has attracted an increasing array of members and conference participants from the broadcasting, information technology, and multimedia sides of our domain. We are benefiting from the growth of this 'productive anarchy.' We have more work to do in this regard, but converging technologies as well as the convergence of policy imperatives will lead to greater personal interaction between the individuals in these spheres of activity. Fortunately, PTC provides a forum for such interaction at the worldwide level through the annual conference, our regional seminars, the PTR (*Pacific Telecommunications Review*), and other activities.

James Savage is the Assistant Director, Pacific Telecommunications Council and the Editor of the Pacific Telecommunications Review.

INS leads the way

TEST INSTRUMENTATION • NETWORK MANAGEMENT SYSTEMS

TDR/LAN SCANNER

- TDR and cable scanning for UTP, STP, Coax, IBM cabling.
- Measures NEXT, length, ATTENUATION, Resistance, continuity, LAN monitor.
- Optional cable management and cable grading software.

INTERVIEW 8000 WAN ANALYSERS

- High performance protocol test system for field and LAB use.
- Supporting C, SS7, Frame Relay, ISDN, X25, SNA, QLLC and user decodes.
- Interfaces for RS232, V.35, X.21, G703, T1.

PC-BASED BERT

- Turn your PC into high speed BERT.
- Support RS232, V.35, X21, G703, G704.
- G821 analysis with histogram and capture to disk.
- G704 monitor tool.

MATRIX SWITCH

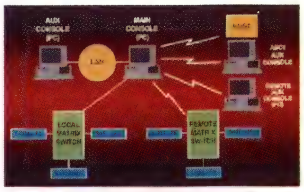

- Non-blocking
- 12 to 4096 ports.
- Mix protocol speeds and interfaces.
- Netview support.

LIGHT SOURCE/OTDR/ POWER METER

- Fiberoptic test equipment for LAN, WAN, FDDI.
- PC-based and hand-held.
- Power attenuation, dual wave length test systems, LED or Laser sources, Talk sets.

NETWORK PROBE

- Multi-function analyser.
- Built-in interfaces for RS232, X.21, V.35, ISDN.
- Multiple protocols support.
- Light weight/low cost.
- AC/DC operated.

LAN/WAN ANALYSER

LAN: Token Ring
ETHERNET

WAN: ISDN
FASTPAC
AUSTPAC
TRANSEND
DDN
SNA

PROTOCOLS:

X.25, SNA, HDLC, SDLC, SS7, Q921/Q931, XNS, TCP/IP, NETBIOS, NETBUI, ISO, APPLETALK, MBI, DECNET, SUNNFS, X-WINDOWS, SNMP.

SIMULATORS

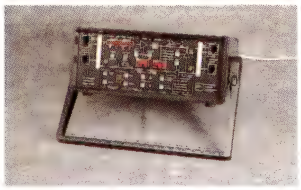

- Most major Telephone network types.
- DDS, ISDN links.
- DDS, ISDN cables.
- Multiple cable gauges, length and speed up to 34.8MBPS.

DATA PATCH/ SWITCH SYSTEMS

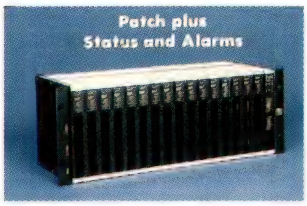

- Circuit switching and PATCHING.
- RS232,. V.35, X.21, G703 etc.
- Transparent to protocol.
- Full remote or local control.

PC-BASED PROTOCOL ANALYSER

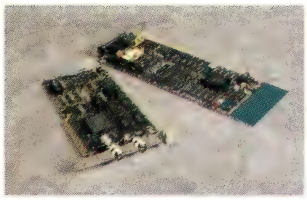

- Turn PC into full function protocol analyser.
- Half size card supporting major protocols, interfaces.
- Speed up to 256KBPS.

SOFTWARE

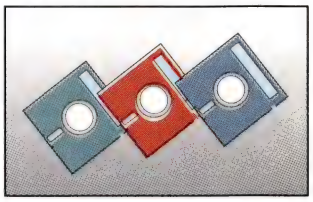

- LAN management software for Ethernet, Token Ring, ARCNET and Appletalk.
- Comprehensive statistics and protocol decodes.
- Low, low, cost!

Please send me more information on:

☐ TDR/LAN SCANNER ☐ INTERVIEW 8000 ☐ PC-BASED BERT ☐ MATRIX SWITCH ☐ OTHERS
☐ NETWORK PROBE ☐ LAN/WAN ANALYSER ☐ SIMULATOR ☐ DATA PATCH/SWITCH SYSTEM
☐ FIBEROPTIC TEST EQUIPMENT ☐ SOFTWARE ☐ PC-BASED PROTOCOL ANALYSER

NAME: _____ ADDRESS: _____

ORGANISATION: _____ PHONE: _____

INS
ACN. 003 338 476

INFORMATION NETWORK SOLUTIONS PTY LTD
Unit 3a, 3 Lanceley Place Artarmon, NSW, 2064
PHONE: (02) 906 - 6335 FAX: (02) 906 - 6327

ATUG AUSTRALIAN TELECOMMUNICATIONS USERS GROUP

From the Executive Director

Wally Rothwell

Best Practice in Network Management

There is a very real need for today's communications managers to come to a common understanding of what is best network management practice and, if possible, to benchmark it.

As Shell's Ray Steele said at ATUG'93: "I suggest that we [Communications Managers] are soon to be told by our management and our customers that we will benchmark and we will deliver metrics of our own performance and we will be compared. We are all involved to stay within best management practice, the process of re-engineering our business to best suit our customers and our management."

Thus, in order to further that common understanding of network management, ATUG commissioned a study by Strategic Technology Management. The study was fully sponsored by seven vendor and three user organisations.

The five phase study, now in its third phase of analysis and interpretation of 15 major interviews and some 50 to 60 questionnaires, has already brought forth some very useful information. This was presented by one of the two principals of STM, John Hont, at ATUG'93.

The final report will soon be published and made available to ATUG members.

The Interim Report

With the aim of trying to define a role for network management and its future directions, it would be useful also to be able to give a statement of management requirements mutually understood by both users and suppliers. Products currently available deliver much of what is required, but without that common understanding of the requirements, as well as how to use the products, the result often falls short of what the Communications Manager's customers are hoping will be delivered.

The situation at present is that, broadly speaking, larger users seem to be following one of two paths. Some are integrating their voice and data networks at the carrier interface, while others have opted for voice VPNs and a private data network. VPNs and LANs are becoming increasingly popular.

There is still no such thing as a pure network, most being a collection of structures, protocols and systems. OSI is spreading steadily, but proprietary systems still rule and are likely to survive for a few years yet, such is the magnitude of capital investment tied to them.

Consequently, network managers are forced to struggle with a medley of non-communicating NMSs (Network Management Systems). Each has to be used independently. The current trend towards distributed applications and client-server environments accentuates this need for a single integrated network management tool.

Secondly, the investigation has drawn our attention to several significant problems that are currently facing network managers.

IT and communications networks are no different from any other part of business. They have to produce identifiable returns and show that they are efficient. Network managers find that they are being held ever more accountable, but they have no satisfactory way of quantifying network performance and efficiency.

Benchmarking and Billing

ATUG sees this need for a system of benchmarking as one of the main issues that will come out of the study and hopes that the latter will stimulate significant progress in this area. Billing is another problem area. The need for greater internal accountability raises the issue of the importance of a satisfactory billing system to the network manager. Data NMSs are deficient in this area, and there is no NMS that permits internal billing of LAN and WAN traffic. Network managers would very much like to be able to bill users of such facilities.

The situation is further complicated by carrier billing. There is a lack of any standardisation of billing periods for leased facilities, in terms of both line rentals and usage charges. Within one specific network site there might be a number of leased lines, all with different billing intervals.

Add a few Flexi-Plans to taste, with their own unique periods and the network manager is presented with a quite indigestible menu.

Thirdly, there is no NMS for voice available that combines traffic dispersion, fault reporting and route utilisation information across the network. With a still significant reliance on user complaints for fault detection, a more proactive approach would be beneficial. After all, it is on uptime and application response time that users judge their network performance.

So what do users want from an NMS? They would like to manage their networks with one universal network management tool. They do not want a lowest common denominator product, a 'jack of all trades but master of none.' The NMS must be able to provide business oriented reporting and operational performance indicators for both specific sectors and the whole network.

Histograms could be used to show, at a glance, summaries of performance on a daily or longer term basis. Users realise that no single vendor is going to have the best tool for every system, and appreciate that they will have to continue to employ a number of systems. However, they must integrate, so that a single tool can present an overview.

In addition to the essential benchmarking issue, ATUG sees the need for close communication between the vendor and user as an issue worthy of further major attention. We hope that this study, when complete, will enable vendors to better understand what users want out of an NMS, thus ensuring that vendor offerings are developed according to the users' needs.

Wally Rothwell
Executive Director

Checklist Solution For Your Carrier's Performance

ATUG has always been proactive in the development of Australian standards, through its representation on Austel's Standards Advisory Committee and Austel standards working groups, Standards Australia technical committees and CCITT committees. The aim being cost-effective and user-friendly standards and, where possible, adoption or adaptation of international standards. Standards allow for a level playing field where the best and the worst can be measured. In most cases, the differences are physical and there is no confusion between the best and the worst, but in the case of telephone calls the success of a call or the failure of that call is mostly judged by whether the call was clear and uninterrupted.

Significant Milestone

However, the success or failure of a call is crucial when it comes to emergency situations and business transactions. In Australia, thus far, being able to rate the success of a call or the carrier's performance generally, based on agreed measures, has never been attempted. Nevertheless, the need has always been there.

A draft checklist of carrier performance measures has recently been devised for the first time in Australia by Standards Working Group 12/1 which is involved with end-to-end network service performance standards. According to ATUG director and winner of the 1993 Excellence in Communications Management award, Allan Horsley, this checklist is a significant milestone for users. It includes aspects such as time to dial tone, post dialling delay, inaccurate call connection, call drop-out, distortion and loss factors, crosstalk and call clearing delay. Similar measurable performance standards are used in New York in the decision process for carrier licence renewals.

The working group has developed 20 telephone service performance parameters as a guide for users and carriers so that artificial user expectations are eliminated but, at the same time, act as objectives for carriers to achieve.

The 20 draft parameters are classified as measures concerned with speed, accuracy and dependability. They are also used to rate call connection set-up, user information transfer and call connection disengagement release.

ATUG Action

ATUG requests its members to advise the working group on this draft checklist of carrier performance measures by writing or faxing comments.

Letters should be addressed to: Carrier Performance Parameters, ATUG, P.O. Box 357, Milsons Point, NSW, 2061. Facsimiles may be forwarded to ATUG on (02) 925 0880.

1. TIME TO DIAL TONE

This parameter measures the time interval between lifting the handset and reception of dial tone. According to the working group, 98% of calls should receive a dial tone within three seconds.

2. POST DIALLING DELAY

This is the time interval for a ring to be received between the end of user or terminal equipment dialling and the reception of an appropriate network response. The working group has not decided on the appropriate time limit

Telephone Service Performance Parameters

Criterion/Function	Speed	Accuracy	Dependability
Connection Set-Up	1. Time to Dial Tone 2. Post Dialling Delay	3. Probability of Inaccurate Call Connection	4. Probability of Connection Set-Up Failure
User Information Transfer	5. User Information Transfer Delay	6. Atten./Freq. Distortion 7. Intermodulation Distortion 8. Harmonic Distortion 9. Amplitude Clipping 10. Continuous Random Noise 11. Impulse Noise 12. Speech Clipping 13. Group Delay Distortion 14. Echo Loss 15. Crosstalk 16. Transmission Loss	17. Probability of Call Cut-Off/Interruption
Connection Disengagement Release	18. Call clearing Delay	19. Probability of Inaccurate Call Clearing	20. Probability of Call Clearing Failure

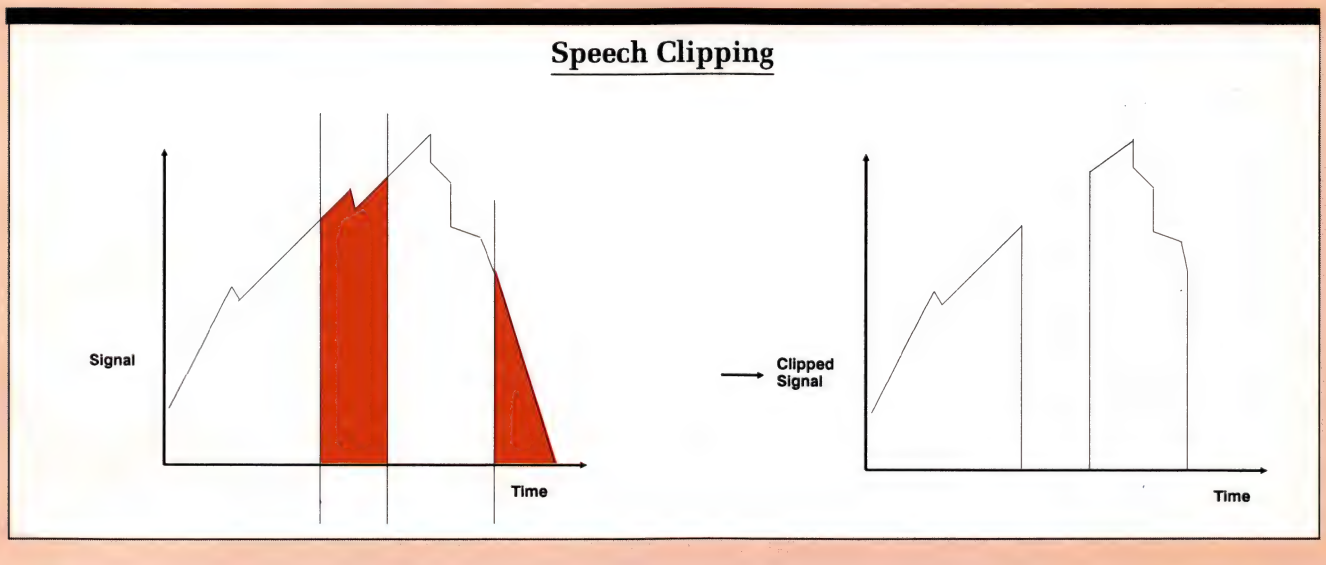

but a certain percentage of calls should receive a ring within a few seconds. Details of this measure have yet to be confirmed.

3. PROBABILITY OF INACCURATE CALL CONNECTION

The working group has not finalised the values for this measure but it believes that the probability of inaccurate call connection should equal the number of inaccurate call connections divided by the total number of connection access attempts. This measure aims to find out whether incorrect call set up is due to:

a) Preselection not working but the call is connected to the desired party;
b) Charging function inaccuracies, for example, incorrect timing, multiple meter pulses; or
c) From the wrong calling line identification message.

4. PROBABILITY OF CONNECTION SET UP FAILURE

This parameter may be used to find the probability of a valid bid for service resulting in one of the following conditions:

a) Dial tone returned after dialling completed;
b) No ring and no answer;
c) All circuits busy signal or announcement;
d) Connection to the wrong number (misrouting); and
e) Double connection.

Again, the working group has not decided on the exact value of this parameter but believes the value should equal the probability of connection set-up failure which is the number of connection set-up failures divided by the total number of connection access attempts. Based on traffic design information this figure may be below 0.1 per cent or below two per cent based on CCITT Rec E845.

5. USER INFORMATION TRANSFER DELAY

This parameter measures the time taken for a signal to travel from one end of an established connection to the other. According to the working group, this parameter should be split into two categories to cater for terrestrial circuits and satellite circuits, where the delay for terrestrial circuits is less than or equal to 200ms, and for satellite circuits user information transfer delay is less than or equal to 470ms.

6. ATTENUATION/FREQUENCY DISTORTION

This parameter measures the difference between the actual response of voltage versus frequency and the ideal (planned) response of voltage versus frequency, referred to as the corresponding difference at 1000Hz. However, more work is required to assess these values.

7. INTERMODULATION DISTORTION

This is the occurrence of the sum and difference frequencies due to non-linearities in the characteristic of the transmission system resulting in the distortion of the original waveform. More work is required to assess these values.

8. HARMONIC DISTORTION

This parameter measures the occurrence of additional frequencies that are whole number multiples of the original frequencies. This is also due to non-linearities in the transmission system resulting in the distortion of the original waveform. Further work is planned to determine the value.

9. AMPLITUDE CLIPPING

Amplitude clipping is the limitation im-

SYDNEY
Level 11, 80 Alfred Street
Milsons Point, NSW, 2061
Australia
Tel: (02) 957 1333, (008) 22 6281
Fax: (02) 925 0880
MELBOURNE
Unit 3, 8-12 Sandilands Street
South Melbourne, VIC, 3205
Australia
Tel: (03) 690 6395
Fax: (03) 699 9094

posed on the maximum amplitude of a signal in a transmission system. This is due to the inability of the system to accept the full excursion of the amplitude of the signal and results in the distortion of the original waveform. The working group will examine this further.

10. CONTINUOUS RANDOM NOISE

Continuous random noise is any unwanted signal in an established connection, which consists of a large number of transient disturbances with a statistically random distribution. The working group has considered two figures for this level, one from Telecom and another from TS012. The Telecom level is less than -59.5 dBmOp whilst the TS012 continuous noise level measured at the customer end is quoted as not exceeding -46 dBmOp for more than 20 one minute intervals in any four week period.

11. IMPULSE NOISE

Impulse noise is any unwanted signal that occurs in an established connection, which consists of a non-overlapping random succession of transient disturbances resulting from such sources as switching centres, lightning discharges and dial pulses. The working group has examined figures supplied by Telecom and TS012, and has determined that the impulsive noise counts should not exceed 18 counts above -23 dBmO in a 15 minute period.

SYDNEY
Level 11, 80 Alfred Street
Milsons Point, NSW, 2061
Australia
Tel: (02) 957 1333, (008) 22 6281
Fax: (02) 925 0880
MELBOURNE
Unit 3, 8-12 Sandilands Street
South Melbourne, VIC, 3205
Australia
Tel: (03) 690 6395
Fax: (03) 699 9094

12. SPEECH CLIPPING

This parameter measures the 'cut off' or 'clipping' and complete loss of any part of the signal in time, for example, the loss of a syllable or words in speech. More work is required to assess the values of this parameter.

13. GROUP DELAY DISTORTION

This distortion is the signal due to the variation in time taken for the frequency components of a signal to travel across an established connection. The working group is considering the following options: the highest permissible group delay distortion for the lowest frequency at which the attenuation is less than 10dB below the level at 800Hz to be at 15ms; and the highest frequency at which the attenuation is less than 10dB below the level at 800Hz to be at 7.5ms.

14. ECHO LOSS

This parameter measures the occurrence of delayed signals which occur when more than one image of the same signal is received at different times. Two options are being considered by the working group, one based on TS012, where the echo loss should be greater than or equal to 15dB, and the other based on the G.122 CCITT recommendation.

15. CROSSTALK

Crosstalk is the interference caused to an established connection by the coupling of an unwanted signal with the established connection. The working group believes the level of crosstalk acceptable should not be less than 65dB below the signal in the established connection.

16. TRANSMISSION LOSS

Transmission loss refers to the difference in level between the transmitted signal and the received signal in an established connection. This level may be frequency (bandwidth) dependent. The working group has determined that 90% of end-to-end connections should not have an attenuation of less than 25dB where measurement is based on insertion loss at specific frequency and impedance.

17. PROBABILITY OF CALL CUT OFF

This parameter measures the probability of an established connection being interrupted or completely cut off. 'Interruption' has been defined as an event that temporarily prevents effective communication between users but does not result in disconnection, while 'cut off' is defined as an event that results in a disconnected service. These events cannot be initiated by either user in the established connection. The working group has considered the probability of the premature release for a national circuit should be less than or equal to 2.4 x 10E-4; while the probability of speech loss with duration which is shorter than 10 seconds due to transmission interruptions should be less than 0.5%.

18. CALL CLEARING DELAY

This delay is the time interval between a user replacing the handset and the network being capable of establishing a new connection to that user. The call clearing delay may be different for calling and called parties. According to the working group, 95% of calls should clear down within one second, while 98.5% of calls should clear down within three seconds.

19. PROBABILITY OF INACCURATE CALL CLEARING

This parameter is used to find out the probability of a connection incorrectly disconnecting when the user replaces the handset. In such an instance charging may continue for that call. The probability of incorrect call clearing should be less than or equal to 1/100,000 according to the working group.

20. PROBABILITY OF CALL CLEARING FAILURE

This parameter measures the probability of an established connection not disconnecting upon the user replacing the handset. For example, call clearing failure could result in no network action at all, or partial action or that the call may not clear within the specified time. There is more work required by the working group on this parameter.

Toula Mantis

Frames and the Two Cells: Interworking Frame Relay, Fastpac and ATM

The confusion in the marketplace regarding the technologies of Frame Relay (FR) and Asynchronous Transfer Mode (ATM) has prompted ATUG to work with the Frame Relay Forum in an attempt to explain these technologies. Thus, a series of articles from the Forum will be published in these pages. Currently, the Frame Relay Forum, the ATM Forum and the Switched Multimegabit Data Service (SMDS) Interest Group are preparing a joint white paper that addresses the relative positioning of these technologies.

Let's start where we want to end up. This is the mythical world of 'Network Nirvana.' Communications anywhere, anytime, any type; at affordable costs. Ideally, there should be one network capable of supporting all communications services.

This is the goal of the Fast Packet Switched Broadband-ISDN (B-ISDN) program underway within the International Telegraph and Telephone Consultative Committee (CCITT). This network will replace the present day multiple service, multiple platform environment with a streamlined environment in which a variety of services are provided in a fully integrated manner. B-ISDN will support a limited number of connection types and a limited number of user and network interfaces. Fast packet switching will restrict the current explosion of network services while delivering full featured functionality. This will happen! B-ISDN will be here as an integrated full service offering in about 10 years — take note of this; not 10 months as some vendors and carriers would have us believe.

ATM, as the technology of choice in the CCITT's B-ISDN programme, has the power to seamlessly integrate the local, metropolitan and wide area network arenas, like no other technology before. In time, this will empower functional distributed processing and meaningful multimedia application integration, to the point that the only distinction between network capabilities and capacity for any user will be the cost of getting the packets from source to destination. The barriers will not be technological, but economic. [Networking will never be free. We are at a point analogous to the nuclear power industry in the 1950s when they were predicting that electricity would be so cheap that the utilities would not be able to afford to meter customer's use — the analogy should stop there].

This is the dilemma facing the average network user. How do we invest wisely in today's functional solutions while making a meaningful pre-investment in the solutions coming down the pipe? The white paper from the three Forums will help to address this issue in a way that will be relevant for the average user.

Integrated Approach

At present, all communications services like Fastpac, Narrowband ISDN, X.25, DDN, TDM Muxing and PSTN are provided by means of service-specific platforms. Fast packet switching (both cell and frame relay) technologies are capable of being delivered from integrated platforms although they will not be initially deployed in this manner. Specific FR, Metropolitan Area Network (MAN) and ATM platforms will be the 'norm.'

The Frame Relay Forum is promoting FR as the access mechanism for all switched platforms. Packet overhead constraints restrict the minimum workable speed of FR connections to a lower limit of approximately 64Kbps and to an upper limit of around 45Mbps. Any switching platform that can sustain these speeds can be used for frame relay connectivity. We see international examples of private FR switches and public FR over Fastpac (read SMDS), ISDN and X.25 networks. We believe the deployment of FR in Australia will use a similar mix of platforms.

Internetworking Developments

As B-ISDN develops, we will be able to connect two frame relay users over disparate backbones as well as interwork different facilities, for example, connecting a frame relay device to a Fastpac device or X.25 facility. There are two prime ways of achieving this interworking: *Encapsulation* or *Translation*

Encapsulation works by allowing frame relay frames to be written onto the payload of an ATM cell, for example. The envelope is opened at the destination and the contents revealed. Translation requires the creation of equivalent functions in the communicating devices and the pairing of these functions, in order to extract meaningful information from the packets flowing between the devices.

The translation approach can make more optimum use of the network resources, but encapsulation is inherently easier to establish and will probably happen first. Similar techniques can be used to interwork Fastpac and SMDS services, which are based on Distributed Queue Dual Bus (DQDB) technology. Both techniques require agreed standards to allow interworking between networks.

The ability of FR to operate independently of the switching fabric is seen as one of the real determinants of FR's ability to deliver the investment protection users are seeking.

Further articles in this series will describe the interoperability work programs of the ATM and FR Forums as well as exploring specific implementations relevant to the Australian networking environment.

July

5-7 **The Pan Asia Optical Fibre Summit '93**, Conrad Hotel, Hong Kong. This conference will deal with key issues in the field of optical fibre technology, including regional government policy for optical fibre development; regional carrier infrastructure and service roll-out plans; commercial opportunities in the Asia Pacific region; and new directions in optical fibre terrestrial and submarine technology. There will be Ministerial presentations from China, Indonesia and the Philippines, and key industry speakers from regional carriers and leading telecommunications product vendors. A special feature will be the Chinese Government's discussion of its privatisation plans. Fee: $US1,595. Enquiries — AIC Conferences Tel: +852-520-1481 Fax: +852-866-7340.

7-9 **The 1993 Pan Asian Mobile Communications Summit**, Singapore. This conference will be held in both Singapore and Hong Kong, and will cover mobile services policy and planning in Asian markets, GSM Asia-Pacific roll out experiences, satellite personal communications, Digital AMPs, JDC and CDMA implementation experiences, and regional developments in Australia, Hong Kong, Singapore, Indonesia, Malaysia, Philippines and Japan. The Hong Kong conference will be held on July 12-14. Fee: $US1,895. Enquiries — IIR Conferences Tel: (02) 954 5844 Fax: (02) 959 4684.

12-14 **Network Cabling Design**, Maritime Conference Centre, Sydney. This intensive hands-on seminar is designed for people involved in managing and installing networks, and covers areas such as cabling standards; cable construction, uses and transmission methods; network test equipment; preparing cabling cost estimates; and future cabling trends. The course will also be held in Melbourne later in the year. Fee: $1,445. Enquiries — IIT Training Tel: (02) 959 5990 Fax: (02) 956 6375.

22-23 **Auditing Novell NetWare**, Renaissance Hotel, Sydney. This course is designed for computer auditors who need to audit Novell NetWare, and provides both an overview of the technical fundamentals of NetWare and a practical guide to auditing it. Attendees will learn the exposures and control concepts specific to NetWare, recommended procedures for protecting the integrity of LANs and the specific access control facilities critical to LAN implementation and administration. Participants need a good understanding of PCs, DOS and the DOS batch language. Fee: $1,195. Enquiries — AIC Training Tel: (02) 235 1700 Fax: (02) 223 8216.

August

4-5 **Outsourcing & Systems Integration**, Lakeside Hotel, Canberra. This event is designed for managers whose organisation is examining the potential advantages of outsourcing in order to manage costs more effectively and concentrate on core business processes. Enquiries — AIC Conferences (02) 235 1700 Fax: (02) 223 8216.

4-7 **TUANZ 93**, Ellerslie Function Centre, Auckland. This annual conference and exhibition will explore the issues, applications and converging technologies that will lead the way to the year 2000. The theme of 'Towards 2000' will be addressed from several perspectives: visionary, business, carrier and Pacific. For the first time the exhibition will be open through to Saturday, to enable those operating small businesses to attend. A series of seminars on Saturday, 7 August will be aimed specifically at the educational sector, and will look at the role of telecommunication in distance education, educational TV and other areas of interest. Fee: $NZ750 (Members), $NZ975 (Non-members). ATUG members are entitled to attend at the member rate. Enquiries — TUANZ Tel: +64 9 488 1602 Fax: +64 9 489 9515.

11-13 **Open Systems Interconnection**, Hotel Nikko, Potts Point. This conference takes an in-depth look at the Australian Government GOSIP III and AGGOS policies and the move towards an open systems computing environment. Major Unix vendors will discuss their initiatives in open systems, and speakers from government, commercial organisations and industry bodies will discuss their practical experiences. The second day of the conference will focus on OSI, and an optional one-day workshop will be held on Day Three. This event will also be held in Canberra on August 16-18. Fee: Conference only $1,295; conference and workshop $1,895; workshop only $795. Enquiries — IIR Conferences Tel: (02) 954 5844 Fax: (02) 959 4684.

11-13 **Executive Information Systems**, Hilton Hotel, Sydney. This seminar will explore the emerging opportunities offered by EIS. Eleven leading Australasian organisations will present papers detailing their experiences, and explaining how an EIS can help maintain a competitive edge through strategic management of corporate information. Enquiries — AIC Conferences Tel: (02) 235 1700 Fax: (02) 223 8216.

16-18 **Implementing GSM Networks and Services**, Hotel Intercontinental, Sydney. This conference examines the critical issues surrounding GSM, and features a comprehensive account of global trends. There will be eight international speakers from France, Sweden, Finland, Germany, Luxembourg and the U.K. who will discuss issue such as the GSM services and facilities which will be available to users, the implications for GSM technology in the marketplace, how the GSM networks will interface, international roaming issues, and GSM security requirements. There will also be an optional one-day workshop. Fee: Two-day conference $1,295, One-day workshop $795, conference and workshop $1,895. Enquiries — IIR Conferences Tel: (02) 954 5844 Fax: (02) 959 4684.

23-27 **Interop 93**, Moscone Convention Center, San Francisco, California. The theme of this year's Interop is 'At Work and Working Together,' and the conference and exhibition program promises more than 400 industry-leading network computing and open systems companies. There will be over 110 conference sessions, tutorials and workshops, and attendees can see state-of-the-art solutions and be shown how these new products interoperate with the widely installed technologies of today. Courses from beginner to expert level are offered, and this year there are new one-day training workshops especially designed for those with tight schedules. Enquiries — Interop Company, Tel: +1 415 941 3399 Fax: +1 415 949 1779.

September

1-3 **Asia Pacific Mobile Communications — Implementing GSM & USDC**, Shangri-La Hotel, Singapore. This annual conference will examine the future of mobile communications in the region, specifically focusing on implementing GSM and USDC digital cellular technologies with reports of the latest developments in Europe, USA and Japan. A technical day will concentrate on the practicalities of creating a successful and profitable network. Fee: $S2,350. Enquiries — IBC Technical Services Tel: +65 732 1970.

20-22 **Understanding Data Communications Networks**, Maritime Conference Centre, Sydney. This course aims to give professionals working in data communications a thorough grounding in communications protocols and standards, modem terminology and technology, hardware and software interfaces, multiplexers, and transmission media and their uses. Fee: $1,445. Enquiries — ITT Training Tel: (02) 959 5990 Fax: (02) 956 6375.

27-1/10 **ASWEC '93**, Hyatt Kingsgate, Sydney. The 7th Australian Software Engineering Conference has the theme this year of 'Software Quality and Other Urban Myths,' and will feature papers on CASE, development environments, management of software developments and technical innovations. Enquiries — Institution of Radio and Electronics Engineers Australia Tel: (02) 327 4822 Fax: (02) 362 3229.

October

12-14 **The Inmarsat International Conference and Exhibition on Mobile Satellite Communications**, CNIT, Paris. The first Inmarsat conference and exhibition in 1989 focused on the development of mobile communications. This year the conference will be addressing the future of mobile satellite communications, and senior-level speakers will present their views on a range of issues covering all aspects of the industry. The format will consist of plenary and streamed sessions, and there will also be a series of open forums. The exhibition will show the latest equipment in the field of mobile satellite communications. Enquiries — Tania Starley, IBC Technical Services Tel: +44 71 637 4383 Fax: +44 71 631 3214.

24-27 **The IREE Communications Conference**, Hilton International Hotel, Melbourne. The IREE has announced this new conference, which will replace the former IREECON. This year's inaugural event will have the theme 'Communications: Foundations for the Future,' and will run streams on radiocommunications, telecommunications and industry support. The keynote speaker will be the newly elected Director of the ITU Telecommunications Development Bureau, Mr Arnold Djiwatampu. A highlight of the three day event will be a special 'Hypothetical' session led by John Leaney, Senior Lecturer, University of Technology, Sydney, and featuring a panel of key industry figures. There will also be a trade exhibition of the latest technology and equipment. Enquiries — Tel: (02) 327 4822 Fax: (02) 362 3229.

November

16-18 **Middle East & Gulf Mobile Communications — Implementing New Digital Mobile Communications Systems**, Inter-Continental Hotel, Dubai. Concentrating on GSM and being held at a critical time in the region's development of mobile systems, this conference will examine operational experiences with the implementation of GSM cellular networks, the issues surrounding standards and roaming, and the future impact of other mobile services. A technical day will allow delegates to understand the practicalities of building a successful and profitable network. Enquiries — IBC Technical Services Tel: +44 71 637 4388 Fax: +44 71 631 3214.

28-1/12 **ACOFT-18 '93**, Northbeach Parkroyal Hotel, Wollongong. The 18th Australian Conference on Optical Fibre Technology will cover the latest research, developments, production applications and business strategies of optical fibres, waveguides, sources, detectors and other services for the telecommunications and sensors industries. A trade exhibition will be held in conjunction with the conference. Enquiries — Conference Secretary, IREE Tel: (02) 327 4822 Fax: (02) 362 3229.

To provide communications of the future, Australia's telephone exchanges are being updated now.

Not long ago, the video phone was something you would only see in science-fiction movies.

But before the end of the decade it will be a reality, in thousands of homes and businesses across Australia.

And many more exciting developments in the world of communications will soon become available, thanks to the new Alcatel 1000 System 12 currently being delivered to Australia's telephone exchanges.

Of course, this is an enormous, highly technical and costly task. But it's satisfying to know that one company in Australia has the engineering and manufacturing know-how to build and supply this state of the art equipment.

Alcatel Australia. We're Australia's largest communications company because we know how to compete.

ALCATEL
AUSTRALIA

C&M ALC 0063/R

Chosen by the most discerning Business and Government Institutions

IQ INTERQUARTZ®

The market leaders in Business Phones

IQTEL 500

MORE FEATURES AT A BUDGET PRICE

- Liquid crystal display
- Manual timer
- On hook / monitor dialling with auto release
- 13 memory capacity,10 two touch, 3 one touch, all 16 digit capacity
- Swing out index label
- Hearing aid compatible
- Optional headset support
- Plus all IQTEL common features
- Colour Grey
- Two year warranty

INTERQUARTZ (A'ASIA) PTY. LTD

MELBOURNE
42 Dougharty Road
West Heidelberg VIC 3081
Telephone (03) 457 5000
Fax (03) 458 4049

SYDNEY
8 Pioneer Avenue
Thornleigh NSW 2120
Telephone (02) 875 4011
Fax (02) 875 4146

CANBERRA
Telephone (06) 243 4860
Fax (06) 243 4848

BRISBANE
4/12 Tolmer Place
Spingwood QLD 4127
Telephone (07) 808 8008
Fax (07) 808 8666

TOLL FREE 008 333 131 AUSTRALIA

AUSTEL
AUSTRALIAN TELECOMMUNICATIONS AUTHORITY

Only products that carry an AUSTEL permit have been tested to meet AUSTEL's safety and technical standards